GNU Astronomy Utilities

Astronomical data manipulation and analysis programs and libraries
for version 0.3, 1 June 2017

Mohammad Akhlaghi

Gnuastro (source code and book) authors (sorted by number of commits in project history):

Mohammad Akhlaghi (akhlaghi@gnu.org, 696)
Mosè Giordano (mose@gnu.org, 29)
Vladimir Markelov (vmatroskin@gmail.com, 15)
Boud Roukema (boud@cosmo.torun.pl, 3)

This book documents version 0.3 of the GNU Astronomy Utilities (Gnuastro). Gnuastro provides various programs and libraries for astronomical data manipulation and analysis.

Copyright © 2015-2017 Free Software Foundation, Inc.

> Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

For myself, I am interested in science and in philosophy only because I want to learn something about the riddle of the world in which we live, and the riddle of man's knowledge of that world. And I believe that only a revival of interest in these riddles can save the sciences and philosophy from narrow specialization and from an obscurantist faith in the expert's special skill, and in his personal knowledge and authority; a faith that so well fits our 'post-rationalist' and 'post-critical' age, proudly dedicated to the destruction of the tradition of rational philosophy, and of rational thought itself.

—*Karl Popper. The logic of scientific discovery. 1959.*

Short Contents

1 Introduction 1
2 Tutorials 14
3 Installation 25
4 Common program behavior 48
5 Data containers 80
6 Data manipulation 97
7 Data analysis 148
8 Modeling and fitting 195
9 High-level calculations 216
10 Library 223
11 Developing 320
A Gnuastro programs list 341
B Other useful software 343
C GNU Free Doc. License 346
Index: Macros, structures and functions 354
Index 358

Table of Contents

1 Introduction ... 1
1.1 Quick start ... 1
1.2 Science and its tools ... 2
1.3 Your rights ... 4
1.4 Naming convention ... 5
1.5 Version numbering ... 5
1.5.1 GNU Astronomy Utilities 1.0 ... 6
1.6 New to GNU/Linux? ... 7
1.6.1 Command-line interface ... 7
1.7 Report a bug ... 9
1.8 Suggest new feature ... 11
1.9 Announcements ... 11
1.10 Conventions ... 12
1.11 Acknowledgments ... 12

2 Tutorials ... 14
2.1 Hubble visually checks and classifies his catalog ... 14
2.2 Sufi simulates a detection ... 17

3 Installation ... 25
3.1 Dependencies ... 25
3.1.1 Mandatory dependencies ... 26
3.1.1.1 GNU Scientific library ... 26
3.1.1.2 CFITSIO ... 26
3.1.1.3 WCSLIB ... 28
3.1.2 Optional dependencies ... 28
3.1.3 Bootstrapping dependencies ... 29
3.2 Downloading the source ... 31
3.2.1 Release tarball ... 31
3.2.2 Version controlled source ... 32
3.2.2.1 Bootstrapping ... 33
3.2.2.2 Synchronizing ... 35
3.3 Build and install ... 36
3.3.1 Configuring ... 36
3.3.1.1 Gnuastro configure options ... 37
3.3.1.2 Installation directory ... 39
3.3.1.3 Executable names ... 43
3.3.1.4 Configure and build in RAM ... 44
3.3.2 Tests ... 45
3.3.3 A4 print book ... 45
3.3.4 Known issues ... 46

4 Common program behavior 48
4.1 Command-line 48
4.1.1 Arguments and options 48
4.1.1.1 Arguments 49
4.1.1.2 Options 50
4.1.2 Common options 52
4.1.2.1 Input/Output options 52
4.1.2.2 Processing options 54
4.1.2.3 Operating mode options 55
4.2 Configuration files 59
4.2.1 Configuration file format 59
4.2.2 Configuration file precedence 60
4.2.3 Current directory and User wide 61
4.2.4 System wide 61
4.3 Multi-threaded operations 62
4.3.1 A note on threads 62
4.3.2 How to run simultaneous operations 63
4.4 Numeric data types 64
4.5 Tables 66
4.5.1 Recognized table formats 67
4.5.2 Gnuastro text table format 69
4.5.3 Selecting table columns 71
4.6 Tessellation 72
4.7 Getting help 74
4.7.1 `--usage` 74
4.7.2 `--help` 75
4.7.3 Man pages 76
4.7.4 Info 76
4.7.5 help-gnuastro mailing list 77
4.8 Automatic output 77
4.9 Output headers 78

5 Data containers 80
5.1 Fits 80
5.1.1 Invoking Fits 82
5.1.1.1 HDU manipulation 83
5.1.1.2 Keyword manipulation 84
5.2 ConvertType 87
5.2.1 Recognized file formats 87
5.2.2 Color 89
5.2.3 Invoking ConvertType 90
5.3 Table 94
5.3.1 Invoking Table 94

6 Data manipulation 97
6.1 Crop 97
6.1.1 Crop modes 97
6.1.2 Crop section syntax 100
6.1.3 Blank pixels 100
6.1.4 Invoking Crop 101
6.1.4.1 Crop options 102
6.1.4.2 Crop output 107
6.2 Arithmetic 108
6.2.1 Reverse polish notation 108
6.2.2 Arithmetic operators 109
6.2.3 Invoking Arithmetic 114
6.3 Convolve 117
6.3.1 Spatial domain convolution 118
6.3.1.1 Convolution process 118
6.3.1.2 Edges in the spatial domain 119
6.3.2 Frequency domain and Fourier operations 120
6.3.2.1 Fourier series historical background 120
6.3.2.2 Circles and the complex plane 122
6.3.2.3 Fourier series 123
6.3.2.4 Fourier transform 125
6.3.2.5 Dirac delta and comb 126
6.3.2.6 Convolution theorem 127
6.3.2.7 Sampling theorem 129
6.3.2.8 Discrete Fourier transform 132
6.3.2.9 Fourier operations in two dimensions 133
6.3.2.10 Edges in the frequency domain 134
6.3.3 Spatial vs. Frequency domain 135
6.3.4 Convolution kernel 135
6.3.5 Invoking Convolve 136
6.4 Warp 138
6.4.1 Warping basics 139
6.4.2 Merging multiple warpings 142
6.4.3 Resampling 142
6.4.4 Invoking Warp 143

7 Data analysis 148
7.1 Statistics 148
7.1.1 Histogram and Cumulative Frequency Plot 148
7.1.2 Sigma clipping 149
7.1.3 Sky value 150
7.1.3.1 Sky value definition 151
7.1.3.2 Sky value misconceptions 152
7.1.3.3 Quantifying signal in a tile 153
7.1.4 Invoking Statistics 154
7.2 NoiseChisel 163
7.2.1 NoiseChisel changes after publication 164
7.2.2 Invoking NoiseChisel 165

7.2.2.1 General NoiseChisel options ... 166
7.2.2.2 Detection options ... 169
7.2.2.3 Segmentation options ... 173
7.2.2.4 NoiseChisel output ... 175
7.3 MakeCatalog ... 175
7.3.1 Detection and catalog production ... 176
7.3.2 Quantifying measurement limits ... 177
7.3.3 Measuring elliptical parameters ... 181
7.3.4 Adding new columns to MakeCatalog ... 184
7.3.5 Invoking MakeCatalog ... 185
7.3.5.1 MakeCatalog input files ... 187
7.3.5.2 MakeCatalog general settings ... 188
7.3.5.3 Upper-limit magnitude settings ... 189
7.3.5.4 MakeCatalog output columns ... 190

8 Modeling and fitting ... 195

8.1 MakeProfiles ... 195
8.1.1 Modeling basics ... 195
8.1.1.1 Defining an ellipse ... 195
8.1.1.2 Point Spread Function ... 196
8.1.1.3 Stars ... 198
8.1.1.4 Galaxies ... 198
8.1.1.5 Sampling from a function ... 199
8.1.1.6 Oversampling ... 200
8.1.2 If convolving afterwards ... 200
8.1.3 Flux Brightness and magnitude ... 201
8.1.4 Profile magnitude ... 201
8.1.5 Invoking MakeProfiles ... 202
8.1.5.1 MakeProfiles catalog ... 203
8.1.5.2 MakeProfiles options ... 203
8.1.5.3 MakeProfiles output ... 210
8.2 MakeNoise ... 210
8.2.1 Noise basics ... 210
8.2.1.1 Photon counting noise ... 210
8.2.1.2 Instrumental noise ... 212
8.2.1.3 Final noised pixel value ... 212
8.2.1.4 Generating random numbers ... 212
8.2.2 Invoking MakeNoise ... 214

9 High-level calculations ... 216

9.1 CosmicCalculator ... 216
9.1.1 Distance on a 2D curved space ... 216
9.1.2 Extending distance concepts to 3D ... 220
9.1.3 Invoking CosmicCalculator ... 221

10 Library 223
10.1 Review of library fundamentals 223
10.1.1 Headers 224
10.1.2 Linking 227
10.1.3 Summary and example on libraries 229
10.2 BuildProgram 230
10.2.1 Invoking BuildProgram 231
10.3 Gnuastro library 233
10.3.1 Configuration information (`config.h`) 234
10.3.2 Multithreaded programming (`threads.h`) 235
10.3.2.1 Implementation of `pthread_barrier` 236
10.3.2.2 Gnuastro's thread related functions 237
10.3.3 Library data types (`type.h`) 238
10.3.4 Library blank values (`blank.h`) 243
10.3.5 Data container (`data.h`) 245
10.3.5.1 Generic data container (`gal_data_t`) 245
10.3.5.2 Dataset size and allocation 250
10.3.5.3 Arrays of datasets 252
10.3.5.4 Copying datasets 252
10.3.6 Dimensions (`dimension.h`) 253
10.3.7 Linked lists (`list.h`) 255
10.3.7.1 List of strings 257
10.3.7.2 List of `int32_t` 258
10.3.7.3 List of `size_t` 259
10.3.7.4 List of `float` 261
10.3.7.5 List of `double` 262
10.3.7.6 List of `void *` 264
10.3.7.7 Ordered list of `size_t` 265
10.3.7.8 Doubly linked ordered list of `size_t` 265
10.3.7.9 List of `gal_data_t` 267
10.3.8 FITS files (`fits.h`) 268
10.3.8.1 FITS Macros, errors and filenames 268
10.3.8.2 CFITSIO and Gnuastro types 269
10.3.8.3 FITS HDUs 269
10.3.8.4 FITS header keywords 270
10.3.8.5 FITS arrays (images) 274
10.3.8.6 FITS tables 275
10.3.9 World Coordinate System (`wcs.h`) 277
10.3.10 Text files (`txt.h`) 279
10.3.11 Table input output (`table.h`) 281
10.3.12 Arithmetic on datasets (`arithmetic.h`) 284
10.3.13 Tessellation library (`tile.h`) 288
10.3.13.1 Independent tiles 289
10.3.13.2 Tile grid 294
10.3.14 Bounding box (`box.h`) 298
10.3.15 Polygons (`polygon.h`) 298
10.3.16 Qsort functions (`qsort.h`) 300
10.3.17 Permutations (`permutation.h`) 301

10.3.18 Statistical operations (statistics.h) 302
10.3.19 Binary datasets (binary.h) 307
10.3.20 Convolution functions (convolve.h) 309
10.3.21 Interpolation (interpolate.h) 310
10.3.22 Git wrappers (git.h) 311
10.4 Library demo programs 312
10.4.1 Library demo - reading a FITS image 312
10.4.2 Library demo - inspecting neighbors 313
10.4.3 Library demo - multi-threaded operation 314
10.4.4 Library demo - reading and writing table columns 317

11 Developing 320
11.1 Why C programming language? 320
11.2 Program design philosophy 322
11.3 Coding conventions 323
11.4 Program source 326
11.4.1 Mandatory source code files 327
11.4.2 The TEMPLATE program 329
11.5 Documentation 330
11.6 Building and debugging 331
11.7 Test scripts 332
11.8 Developer's checklist 333
11.9 Gnuastro project webpage 333
11.10 Developing mailing lists 335
11.11 Contributing to Gnuastro 335
11.11.1 Copyright assignment 336
11.11.2 Commit guidelines 337
11.11.3 Production workflow 338
11.11.4 Forking tutorial 339

Appendix A Gnuastro programs list 341

Appendix B Other useful software 343
B.1 SAO ds9 343
B.1.1 Viewing multiextension FITS images 343
B.2 PGPLOT 344

Appendix C GNU Free Doc. License 346

Index: Macros, structures and functions 354

Index 358

1 Introduction

GNU Astronomy Utilities (Gnuastro) is an official GNU package consisting of separate programs and libraries for the manipulation and analysis of astronomical data. All the programs share the same basic command-line user interface for the comfort of both the users and developers. Gnuastro is written to comply fully with the GNU coding standards so it integrates finely with the GNU/Linux operating system. This also enables astronomers to expect a fully familiar experience in the source code, building, installing and command-line user interaction that they have seen in all the other GNU software that they use. The official and always up to date version of this book (or manual) is freely available under Appendix C [GNU Free Doc. License], page 346, in various formats (pdf, html, plain text, info, and as its Texinfo source) at `http://www.gnu.org/software/gnuastro/manual/`.

For users who are new to the GNU/Linux environment, unless otherwise specified most of the topics in Chapter 3 [Installation], page 25, and Chapter 4 [Common program behavior], page 48, are common to all GNU software, for example installation, managing command-line options or getting help (also see Section 1.6 [New to GNU/Linux?], page 7). So if you are new to this empowering environment, we encourage you to go through these chapters carefully. They can be a starting point from which you can continue to learn more from each program's own manual and fully benefit from and enjoy this wonderful environment. Gnuastro also comes with a large set of libraries, so you can write your own programs using Gnuastro's building blocks, see Section 10.1 [Review of library fundamentals], page 223, for an introduction.

Finally it must be mentioned that in Gnuastro, no change to any program will be released before it has been fully documented in this here first. As discussed in Section 1.2 [Science and its tools], page 2, this is the founding basis of the Gnuastro.

1.1 Quick start

Gnuastro has three mandatory dependencies and three optional dependencies for extra functionality, see Section 3.1 [Dependencies], page 25. The latest official release tarball is always available as `gnuastro-latest.tar.gz` (`http://ftp.gnu.org/gnu/gnuastro/gnuastro-latest.tar.gz`). For better compression (faster download), and robust archival features, an Lzip (`http://www.nongnu.org/lzip/lzip.html`) compressed tarball is also available at `gnuastro-latest.tar.lz` (`http://ftp.gnu.org/gnu/gnuastro/gnuastro-latest.tar.lz`), see Section 3.2.1 [Release tarball], page 31, for more details on the tarball release. If you have downloaded the tarball in the `TOPGNUASTRO` directory and the dependencies are installed, you can unpack, compile, check and install Gnuastro with the following commands. If you use GNU Tar, the same command (`$ tar xf`) can also be used to unpack `.tar.lz` tarballs (the Lzip must already be installed).

```
$ cd TOPGNUASTRO
$ tar xf gnuastro-latest.tar.gz    # This works on '.tar.lz' too.
$ cd gnuastro-X.X                  # Replace X.X with version number.
$ ./configure
$ make -j8                         # Replace 8 with no. CPU threads.
$ make check
$ sudo make install
```

See Section 3.3.4 [Known issues], page 46, if you confront any complications. For each program there is an 'Invoke ProgramName' sub-section in this book which explains how the programs should be run on the command-line. You can read it on the command-line by running the command `$ info astprogname`, see Section 1.4 [Naming convention], page 5, and Section 4.7 [Getting help], page 74. The 'Invoke ProgramName' sub-section starts with a few examples of each program and goes on to explain the invocation details. In Chapter 2 [Tutorials], page 14, some real life examples of how these programs might be used are given.

1.2 Science and its tools

History of science indicates that there are always inevitably unseen faults, hidden assumptions, simplifications and approximations in all our theoretical models, data acquisition and analysis techniques. It is precisely these that will ultimately allow future generations to advance the existing experimental and theoretical knowledge through their new solutions and corrections.

In the past, scientists would gather data and process them individually to achieve an analysis thus having a much more intricate knowledge of the data and analysis. The theoretical models also required little (if any) simulations to compare with the data. Today both methods are becoming increasingly more dependent on pre-written software. Scientists are dissociating themselves from the intricacies of reducing raw observational data in experimentation or from bringing the theoretical models to life in simulations. These 'intricacies' are precisely those unseen faults, hidden assumptions, simplifications and approximations that define scientific progress.

> Unfortunately, most persons who have recourse to a computer for statistical analysis of data are not much interested either in computer programming or in statistical method, being primarily concerned with their own proper business. Hence the common use of library programs and various statistical packages. ... It's time that was changed.
>
> —*F. J. Anscombe. The American Statistician, Vol. 27, No. 1. 1973*

Anscombe's quartet[1] demonstrates how four data sets with widely different shapes (when plotted) give nearly identical output from standard regression techniques. Anscombe argues that "Good statistical analysis is not a purely routine matter, and generally calls for more than one pass through the computer". Anscombe's quartet can be generalized to say that users of a software cannot claim to understand how it works only based on the experience they have gained by frequently using it. This kind of subjective experience is prone to very serious mis-understandings about what it really does behind the scenes and can be misleading. This attitude is further encouraged through non-free software[2]. This approach to scientific software only helps in producing dogmas and an "obscurantist faith in the expert's special skill, and in his personal knowledge and authority"[3].

> Program or be programmed. Choose the former, and you gain access to the control panel of civilization. Choose the latter, and it could be the last real choice you get to make.

---

[1] `http://en.wikipedia.org/wiki/Anscombe%27s_quartet`

[2] `https://www.gnu.org/philosophy/free-sw.html`

[3] Karl Popper. The logic of scientific discovery. 1959. Larger quote is given at the start of the PDF (for print) version of this book.

—*Douglas Rushkoff. Program or be programmed, O/R Books (2010).*

It is obviously impractical for any one human being to gain the intricate knowledge explained above for every step of an analysis. On the other hand, scientific data can be very large and numerous, for example images produced by telescopes in astronomy. This requires very efficient algorithms. To make things worse, natural scientists have generally not been trained in the advanced software techniques, paradigms and architecture that are taught in computer science or engineering courses and thus used in most software. The GNU Astronomy Utilities are an effort to tackle this issue.

Gnuastro is not just a software, this book is as important to the idea behind Gnuastro as the source code (software). This book has tried to learn from the success of the "Numerical Recipes" book in educating those who are not software engineers and computer scientists but still heavy users of computational algorithms, like astronomers. There are two major differences: the code and the explanations are segregated: the code is moved within the actual Gnuastro software source code and the underlying explanations are given here. In the source code every non-trivial step is heavily commented and correlated with this book, it follows the same logic of this book, and all the programs follow a similar internal data, function and file structure, see Section 11.4 [Program source], page 326. Complementing the code, this book focuses on thoroughly explaining the concepts behind those codes (history, mathematics, science, software and usage advise when necessary) along with detailed instructions on how to run the programs. At the expense of frustrating "professionals" or "experts", this book and the comments in the code also intentionally avoid jargon and abbreviations. The source code and this book are thus intimately linked, and when considered as a single entity can be thought of as a real (an actual software accompanying the algorithms) "Numerical Recipes" for astronomy.

The other major and arguably more important difference is that "Numerical Recipes" does not allow you to distribute any code that you have learned from it. So while it empowers the privileged individual who has access to it, it exacerbates social ignorance. For example it does not allow you to release your software's source code if you have used their codes, you can only publicly release binaries (a black box) to the community. Exactly at the opposite end of the spectrum, Gnuastro's source code is released under the GNU general public license (GPL) and this book is released under the GNU free documentation license. You are therefore free to distribute any software you create using parts of Gnuastro's source code or text, or figures from this book, see Section 1.3 [Your rights], page 4. While developing the source code and this book together, the developers of Gnuastro aim to impose the minimum requirements on you (in computer science, engineering and even the mathematics behind the tools) to understand and modify any step of Gnuastro if you feel the need to do so, see Section 11.1 [Why C programming language?], page 320, and Section 11.2 [Program design philosophy], page 322.

Imagine if Galileo did not have the technical knowledge to build a telescope. Astronomical objects could not be seen with the Dutch military design of the telescope. In the beginning of his "The Sidereal Messenger" (1610) he cautions the readers on this issue and instructs them on how to build a suitable instrument: without a detailed description of "how" he made his observations, no one would believe him. The same is true today, science cannot progress with a black box. Before he actually saw the moons of Jupiter, the

mountains on the Moon or the crescent of Venus, he was "evasive" to Kepler[4]. Science is not independent of its tools.

Bjarne Stroustrup (creator of the C++ language) says: "Without understanding software, you are reduced to believing in magic". Ken Thomson (the designer or the Unix operating system) says "I abhor a system designed for the 'user' if that word is a coded pejorative meaning 'stupid and unsophisticated'." Certainly no scientist (user of a scientific software) would want to be considered a believer in magic, or 'stupid and unsophisticated'. However, this can happen when scientists get too distant from the raw data and are mainly indulging themselves in their own high-level (abstract) models (creations). For example, roughly 5 years before special relativity and about two decades before quantum mechanics fundamentally changed Physics, Kelvin is quoted as saying:

> There is nothing new to be discovered in physics now. All that remains is more and more precise measurement.
>
> —*William Thomson (Lord Kelvin), 1900*

A few years earlier, in a speech Albert. A. Michelson said:

> The more important fundamental laws and facts of physical science have all been discovered, and these are now so firmly established that the possibility of their ever being supplanted in consequence of new discoveries is exceedingly remote.... Our future discoveries must be looked for in the sixth place of decimals.
>
> —*Albert. A. Michelson, dedication of Ryerson Physics Lab, U. Chicago 1894*

If scientists are considered to be more than mere "puzzle solvers"[5] (simply adding to the decimals of known values or observing a feature in 10, 100, or 100000 more galaxies or stars, as Kelvin and Michelson clearly believed), they cannot just passively sit back and uncritically repeat the previous (observational or theoretical) methods/tools on new data. Today there is a wealth of raw telescope images ready (mostly for free) at the finger tips of anyone who is interested with a fast enough internet connection to download them. The only thing lacking is new ways to analyze this data and dig out the treasure that is lying hidden in them to existing methods and techniques.

> New data that we insist on analyzing in terms of old ideas (that is, old models which are not questioned) cannot lead us out of the old ideas. However many data we record and analyze, we may just keep repeating the same old errors, missing the same crucially important things that the experiment was competent to find.
>
> —*Jaynes, Probability theory, the logic of science. Cambridge U. Press (2003).*

1.3 Your rights

The paragraphs below, in this section, belong to the GNU Texinfo[6] manual and are not written by us! The name "Texinfo" is just changed to "GNU Astronomy Utilities" or "Gnuastro" because they are released under the same licenses and it is beautifully written to inform you of your rights.

[4] Galileo G. (Translated by Maurice A. Finocchiaro). *The essential Galileo*. Hackett publishing company, first edition, 2008.

[5] Thomas S. Kuhn. *The Structure of Scientific Revolutions*, University of Chicago Press, 1962.

[6] Texinfo is the GNU documentation system. It is used to create this book in all the various formats.

GNU Astronomy Utilities is "free software"; this means that everyone is free to use it and free to redistribute it on certain conditions. Gnuastro is not in the public domain; it is copyrighted and there are restrictions on its distribution, but these restrictions are designed to permit everything that a good cooperating citizen would want to do. What is not allowed is to try to prevent others from further sharing any version of Gnuastro that they might get from you.

Specifically, we want to make sure that you have the right to give away copies of the programs that relate to Gnuastro, that you receive the source code or else can get it if you want it, that you can change these programs or use pieces of them in new free programs, and that you know you can do these things.

To make sure that everyone has such rights, we have to forbid you to deprive anyone else of these rights. For example, if you distribute copies of the Gnuastro related programs, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must tell them their rights.

Also, for our own protection, we must make certain that everyone finds out that there is no warranty for the programs that relate to Gnuastro. If these programs are modified by someone else and passed on, we want their recipients to know that what they have is not what we distributed, so that any problems introduced by others will not reflect on our reputation.

The precise conditions of the licenses for the programs currently being distributed that relate to Gnuastro are found in the GNU General Public license (`http://www.gnu.org/copyleft/gpl.html`) that accompany them. This book is covered by the GNU Free Documentation License (`http://www.gnu.org/copyleft/fdl.html`).

1.4 Naming convention

Gnuastro is a package of independent programs and a collection of libraries, here we are mainly concerned with the programs. Each program has an official name which consists of one or two words, describing what they do. The latter are printed with no space, for example NoiseChisel or Crop. On the command-line, you can run them with their executable names which start with an `ast` and might be an abbreviation of the official name, for example `astnoisechisel` or `astcrop`, see Section 3.3.1.3 [Executable names], page 43.

We will use "ProgramName" for a generic official program name and `astprogname` for a generic executable name. In this book, the programs are classified based on what they do and thoroughly explained. An alphabetical list of the programs that are installed on your system with this installation are given in Appendix A [Gnuastro programs list], page 341. That list also contains the executable names and version numbers along with a one line description.

1.5 Version numbering

Gnuastro can have two formats of version numbers, for official and unofficial releases. Official Gnuastro releases are announced on the `info-gnuastro` mailing list, they have a version control tag in Gnuastro's development history, and their version numbers are formatted like "`A.B`". `A` is a major version number, marking a significant planned achievement (for example see Section 1.5.1 [GNU Astronomy Utilities 1.0], page 6), while `B` is a minor version

number, see below for more on the distinction. Note that the numbers are not decimals, so version 2.34 is much more recent than version 2.5, which is not equal to 2.50.

Gnuastro also allows a unique version number for unofficial releases. Unofficial releases can mark any point in Gnuastro's development history. This is done to allow astronomers to easily use any point in the version controlled history for their data-analysis and research publication. See Section 3.2.2 [Version controlled source], page 32, for a complete introduction. This section is not just for developers and is very straightforward, so please have a look if you are interested in the cutting-edge. This unofficial version number is a meaningful and easy to read string of characters, unique to that particular point of history. With this feature, users can easily stay up to date with the most recent bug fixes and additions that are committed between official releases.

The unofficial version number is formatted like: `A.B.C-D`. `A` and `B` are the most recent official version number. `C` is the number of commits that have been made after version `A.B`. `D` is the first 4 or 5 characters of the commit hash number[7]. Therefore, the unofficial version number '`3.92.8-29c8`', corresponds to the 8th commit after the official version `3.92` and its commit hash begins with `29c8`. The unofficial version number is sort-able (unlike the raw hash) and as shown above is very descriptive of the state of the unofficial release. Of course an official release is preferred for publication (since its tarballs are easily available and it has gone through more tests, making it more stable), so if an official release is announced prior to your publication's final review, please consider updating to the official release.

The major version number is set by a major goal which is defined by the developers and user community before hand, for example see Section 1.5.1 [GNU Astronomy Utilities 1.0], page 6. The incremental work done in minor releases are commonly small steps in achieving the major goal. Therefore, there is no limit on the number of minor releases and the difference between the (hypothetical) versions 2.927 and 3.0 can be a very small (negligible to the user) improvement that finalizes the defined goals.

1.5.1 GNU Astronomy Utilities 1.0

Currently (prior to Gnuastro 1.0), the aim of Gnuastro is to have a complete system for data manipulation and analysis at least similar to IRAF[8]. So an astronomer can take all the standard data analysis steps (starting from raw data to the final reduced product and standard post-reduction tools) with the various programs in Gnuastro.

The maintainers of each camera or detector on a telescope can provide a completely transparent shell script or Makefile to the observer for data analysis. This script can set configuration files for all the required programs to work with that particular camera. The script can then run the proper programs in the proper sequence. The user/observer can easily follow the standard shell script to understand (and modify) each step and the parameters used easily. Bash (or other modern GNU/Linux shell scripts) are very powerful and made for this gluing job. This will simultaneously improve performance and transparency. Shell scripting (or Makefiles) are also very basic constructs that are easy to learn and readily available as part of the Unix-like operating systems. If there is no program to do a desired step, Gnuastro's libraries can be used to build specific programs.

[7] Each point in Gnuastro's history is uniquely identified with a 40 character long hash which is created from its contents and previous history for example: `5b17501d8f29ba3cd610673261e6e2229c846d35`. So the string `D` in the version for this commit could be `5b17`, or `5b175`.

[8] `http://iraf.noao.edu/`

The main factor is that all observatories or projects can freely contribute to Gnuastro and all simultaneously benefit from it (since it doesn't belong to any particular one of them), much like how for-profit organizations (for example RedHat, or Intel and many others) are major contributors to free and open source software for their shared benefit. Gnuastro's copyright has been fully awarded to GNU, so it doesn't belong to any particular astronomer or astronomical facility or project.

1.6 New to GNU/Linux?

Some astronomers initially install and use the GNU/Linux operating systems because the software that their research community use can only be run in this environment, the transition is not necessarily easy. To encourage you in investing the patience and time to make this transition, we define the GNU/Linux system and argue for the command-line interface of scientific software and how it is worth the (apparently steep) learning curve. Section 1.6.1 [Command-line interface], page 7, contains a short overview of the very powerful command-line user interface. Chapter 2 [Tutorials], page 14, is a complete chapter with some real world example applications of Gnuastro making good use of GNU/Linux capabilities written for newcomers to this environment. It is fully explained, easy and (hopefully) entertaining.

You might have already noticed that we are not using the name "Linux", but "GNU/Linux". Please take the time to have a look at the following essays and FAQs for a complete understanding of this very important distinction. In short, the Linux kernel is built using the GNU C library (glibc) and GNU compiler collection (gcc). The Linux kernel software alone is useless, in order have an operating system you need many more packages and the majority of such low-level packages in most distributions are developed as part of the GNU project: "the whole system is basically GNU with Linux loaded". In the form of an analogy: to say "running Linux", is like saying "driving your carburetor".

- `https://www.gnu.org/gnu/gnu-users-never-heard-of-gnu.html`
- `https://www.gnu.org/gnu/linux-and-gnu.html`
- `https://www.gnu.org/gnu/why-gnu-linux.html`
- `https://www.gnu.org/gnu/gnu-linux-faq.html`

1.6.1 Command-line interface

One aspect of Gnuastro that might be a little troubling to new GNU/Linux users is that (at least for the time being) it only has a command-line user interface (CLI). This might be contrary to the mostly graphical user interface (GUI) experience with proprietary operating systems. To a first time user, the command-line does appear much more complicated and adapting to it might not be easy and a little frustrating at first. This is understandable and also experienced by anyone who started using the computer (from childhood) in a graphical user interface. Here we hope to convince you of the unique benefits of this interface which can greatly enhance your productivity while complementing your GUI experience.

Through GNOME 3[9], most GNU/Linux based operating systems now have a very advanced and useful GUI. Since the GUI was created long after the command-line, some wrongly consider the command line to be obsolete. Both interfaces are very useful for different tasks (for example you can't view an image, video, pdf document or web page on

[9] `http://www.gnome.org/`

the command-line!), on the other hand you can't reproduce your results easily in the GUI. Therefore they should not be regarded as rivals but as complementary user interfaces, here we will outline how the CLI can be useful in scientific programs.

You can think of the GUI as a veneer over the CLI to facilitate a small subset of all the possible CLI operations. Each click you do on the GUI, can be thought of as internally running a different CLI command. So asymptotically (if a good designer can design a GUI which is able to show you all the possibilities to click on) the GUI is only as powerful as the command-line. In practice, such graphical designers are very hard to find for every program, so the GUI operations are always a subset of the internal CLI commands. For programs that are only made for the GUI, this results in not including lots of potentially useful operations. It also results in 'interface design' to be a crucially important part of any GUI program. Scientists don't usually have enough resources to hire a graphical designer, also the complexity of the GUI code is far more than CLI code, which is harmful for a scientific software, see Section 1.2 [Science and its tools], page 2.

For programs that have a GUI, one action on the GUI (moving and clicking a mouse, or tapping a touchscreen) might be more efficient and easier than its CLI counterpart (typing the program name and your desired configuration). However, if you have to repeat that same action more than once, the GUI will soon become very frustrating and prone to errors. Unless the designers of a particular program decided to design such a system for a particular GUI action, there is no general way to run any possible series of actions automatically on the GUI.

On the command-line, you can run any series of of actions which can come from various CLI capable programs you have decided your self in any possible permutation with one command[10]. This allows for much more creativity and exact reproducibility that is not possible to a GUI user. For technical and scientific operations, where the same operation (using various programs) has to be done on a large set of data files, this is crucially important. It also allows exact reproducibility which is a foundation principle for scientific results. The most common CLI (which is also known as a shell) in GNU/Linux is GNU Bash, we strongly encourage you to put aside several hours and go through this beautifully explained web page: `https://flossmanuals.net/command-line/`. You don't need to read or even fully understand the whole thing, only a general knowledge of the first few chapters are enough to get you going.

Since the operations in the GUI are very limited and they are visible, reading a manual is not that important in the GUI (most programs don't even have any!). However, to give you the creative power explained above, with a CLI program, it is best if you first read the manual of any program you are using. You don't need to memorize any details, only an understanding of the generalities is needed. Once you start working, there are more easier ways to remember a particular option or operation detail, see Section 4.7 [Getting help], page 74.

To experience the command-line in its full glory and not in the GUI terminal emulator, press the following keys together: `CTRL+ALT+F4`[11] to access the virtual console. To return

[10] By writing a shell script and running it, for example see the tutorials in Chapter 2 [Tutorials], page 14.

[11] Instead of `F4`, you can use any of the keys from `F1` to `F6` for different virtual consoles depending on your GNU/Linux distribution, try them all out. You can also run a separate GUI from within this console if you want to.

back to your GUI, press the same keys above replacing `F4` with `F7` (or `F1`, or `F2`, depending on your GNU/Linux distribution). In the virtual console, the GUI, with all its distracting colors and information, is gone. Enabling you to focus entirely on your actual work.

For operations that use a lot of your system's resources (processing a large number of large astronomical images for example), the virtual console is the place to run them. This is because the GUI is not competing with your research work for your system's RAM and CPU. Since the virtual consoles are completely independent, you can even log out of your GUI environment to give even more of your hardware resources to the programs you are running and thus reduce the operating time.

Since it uses far less system resources, the CLI is also very convenient for remote access to your computer. Using secure shell (SSH) you can log in securely to your system (similar to the virtual console) from anywhere even if the connection speeds are low. There are apps for smart phones and tablets which allow you to do this.

1.7 Report a bug

According to Wikipedia "a software bug is an error, flaw, failure, or fault in a computer program or system that causes it to produce an incorrect or unexpected result, or to behave in unintended ways". So when you see that a program is crashing, not reading your input correctly, giving the wrong results, or not writing your output correctly, you have found a bug. In such cases, it is best if you report the bug to the developers. The programs will also report bugs in known impossible situations (which are caused by something unexpected) and will ask the users to report the bug.

Prior to actually filing a bug report, it is best to search previous reports. The issue might have already been found and even solved. The best place to check if your bug has already been discussed is the bugs tracker on Section 11.9 [Gnuastro project webpage], page 333, at `https://savannah.gnu.org/bugs/?group=gnuastro`. In the top search fields (under "Display Criteria") set the "Open/Closed" drop-down menu to "Any" and choose the respective program or general category of the bug in "Category" and click the "Apply" button. The results colored green have already been solved and the status of those colored in red is shown in the table.

Recently corrected bugs are probably not yet publicly released because they are scheduled for the next Gnuastro stable release. If the bug is solved but not yet released and it is an urgent issue for you, you can get the version controlled source and compile that, see Section 3.2.2 [Version controlled source], page 32.

To solve the issue as readily as possible, please follow the following to guidelines in your bug report. The How to Report Bugs Effectively (`http://www.chiark.greenend.org.uk/~sgtatham/bugs.html`) and How To Ask Questions The Smart Way (`http://catb.org/~esr/faqs/smart-questions.html`) essays also provide some very good generic advice for all software (don't contact their authors for Gnuastro's problems). Mastering the art of giving good bug reports (like asking good questions) can greatly enhance your experience with any free and open source software. So investing the time to read through these essays will greatly reduce your frustration after you see something doesn't work the way you feel it is supposed to for a large range of software, not just Gnuastro.

Be descriptive

Please provide as many details as possible and be very descriptive. Explain what you expected and what the output was: it might be that your expectation was wrong. Also please clearly state which sections of the Gnuastro book (this book), or other references you have studied to understand the problem. This can be useful in correcting the book (adding links to likely places where users will check). But more importantly, it will be very encouraging for the developers, since you are showing how serious you are about the problem and that you have actually put some thought into it. "To be able to ask a question clearly is two-thirds of the way to getting it answered." – John Ruskin (1819-1900).

Individual and independent bug reports

If you have found multiple bugs, please send them as separate (and independent) bugs (as much as possible). This will significantly help us in managing and resolving them sooner.

Reproducible bug reports

If we cannot exactly reproduce your bug, then it is very hard to resolve it. So please send us a Minimal working example[12] along with the description. For example in running a program, please send us the full command-line text and the output with the `-P` option, see Section 4.1.2.3 [Operating mode options], page 55. If it is caused only for a certain input, also send us that input file. In case the input FITS is large, please use Crop to only crop the problematic section and make it as small as possible so it can easily be uploaded and downloaded and not waste the archive's storage, see Section 6.1 [Crop], page 97.

There are generally two ways to inform us of bugs:

- Send a mail to `bug-gnuastro@gnu.org`. Any mail you send to this address will be distributed through the bug-gnuastro mailing list[13]. This is the simplest way to send us bug reports. The developers will then register the bug into the project webpage (next choice) for you.
- Use the Gnuastro project webpage at `https://savannah.gnu.org/projects/gnuastro/`: There are two ways to get to the submission page as listed below. Fill in the form as described below and submit it (see Section 11.9 [Gnuastro project webpage], page 333, for more on the project webpage).
  - Using the top horizontal menu items, immediately under the top page title. Hovering your mouse on "Support" will open a drop-down list. Select "Submit new".
  - In the main body of the page, under the "Communication tools" section, click on "Submit new item".

Once the items have been registered in the mailing list or webpage, the developers will add it to either the "Bug Tracker" or "Task Manager" trackers of the Gnuastro project webpage. These two trackers can only be edited by the Gnuastro project developers, but they can be browsed by anyone, so you can follow the progress on your bug. You are most welcome to join us in developing Gnuastro and fixing the bug you have found maybe a good

---

[12] `http://en.wikipedia.org/wiki/Minimal_Working_Example`

[13] `https://lists.gnu.org/mailman/listinfo/bug-gnuastro`

starting point. Gnuastro is designed to be easy for anyone to develop (see Section 1.2 [Science and its tools], page 2) and there is a full chapter devoted to developing it: Chapter 11 [Developing], page 320.

1.8 Suggest new feature

We would always be very happy to hear of suggested new features. For every program there are already lists of features that we are planning to add. You can see the current list of plans from the Gnuastro project webpage at `https://savannah.gnu.org/projects/gnuastro/` and following "Tasks"→"Browse" on the horizontal menu at the top of the page immediately under the title, see Section 11.9 [Gnuastro project webpage], page 333. If you want to request a feature to an existing program, click on the "Display Criteria" above the list and under "Category", choose that particular program. Under "Category" you can also see the existing suggestions for new programs or other cases like installation, documentation or libraries. Also be sure to set the "Open/Closed" value to "Any".

If the feature you want to suggest is not already listed in the task manager, then follow the steps that are fully described in Section 1.7 [Report a bug], page 9. Please have in mind that the developers are all very busy with their own astronomical research, and implementing existing "task"s to add or resolving bugs. Gnuastro is a volunteer effort and none of the developers are paid for their hard work. So, although we will try our best, please don't not expect that your suggested feature be immediately included (with the next release of Gnuastro).

The best person to apply the exciting new feature you have in mind is you, since you have the motivation and need. In fact Gnuastro is designed for making it as easy as possible for you to hack into it (add new features, change existing ones and so on), see Section 1.2 [Science and its tools], page 2. Please have a look at the chapter devoted to developing (Chapter 11 [Developing], page 320) and start applying your desired feature. Once you have added it, you can use it for your own work and if you feel you want others to benefit from your work, you can request for it to become part of Gnuastro. You can then join the developers and start maintaining your own part of Gnuastro. If you choose to take this path of action please contact us before hand (Section 1.7 [Report a bug], page 9) so we can avoid possible duplicate activities and get interested people in contact.

Gnuastro is a collection of low level programs: As described in Section 11.2 [Program design philosophy], page 322, a founding principle of Gnuastro is that each library or program should be very basic and low-level. High level jobs should be done by running the separate programs or using separate functions in succession through a shell script or calling the libraries by higher level functions, see the examples in Chapter 2 [Tutorials], page 14. So when making the suggestions please consider how your desired job can best be broken into separate steps and modularized.

1.9 Announcements

Gnuastro has a dedicated mailing list for making announcements. Anyone that is interested can subscribe to this mailing list to stay up to date with new releases or when the depen-

dencies (see Section 3.1 [Dependencies], page 25) have been updated. To subscribe to this list, please visit `https://lists.gnu.org/mailman/listinfo/info-gnuastro`.

1.10 Conventions

In this book we have the following conventions:

- All commands that are to be run on the shell (command-line) prompt as the user start with a `$`. In case they must be run as a super-user or system administrator, they will start with a single `#`. If the command is in a separate line and next line `is also in the code type face`, but doesn't have any of the `$` or `#` signs, then it is the output of the command after it is run. As a user, you don't need to type those lines. A line that starts with `##` is just a comment for explaning the command to a human reader and must not be typed.
- If the command becomes larger than the page width a `\` is inserted in the code. If you are typing the code by hand on the command-line, you don't need to use multiple lines or add the extra space characters, so you can omit them. If you want to copy and paste these examples (highly discouraged!) then the `\` should stay.

  The `\` character is a shell escape character which is used commonly to make characters which have special meaning for the shell loose that special place (the shell will not treat them specially if there is a `\` behind them). When it is a last character in a line (the next character is a new-line character) the new-line character looses its meaning an the shell sees it as a simple white-space character, enabling you to use multiple lines to write your commands.

1.11 Acknowledgments

The list of Gnuastro authors is available at the start of this book and the `AUTHORS` file in the source code. Here the authors wish to gratefully acknowledge the help and support they received from other people and institutions who had an indirect (not committed in the version controlled history) role in Gnuastro. The plain text file `THANKS` which is distributed along with the source code also contains this list.

The Japanese Ministry of Science and Technology (MEXT) scholarship for Mohammad Akhlaghi's Masters and PhD period in Tohoku University Astronomical Institute had an instrumental role in the long term learning and planning that made the idea of Gnuastro possible. The very critical view points of Professor Takashi Ichikawa (from Tohoku University) were also instrumental in the initial ideas and creation of Gnuastro. Brandon Invergo, Karl Berry and Richard Stallman also provided very useful suggestions during the GNU evaluation process. Bob Proulx from Savannah, has kindly supported Gnuastro's project webpage on Savannah and the management of its version controlled source server there.

We would also like to gratefully thank Mohammad-reza Khellat, Alan Lefor, Yahya Sefidbakht, Francesco Montanari, Ole Streicher, Nicolas Bouché, Rosa Calvi, Lee Kelvin, Guillaume Mahler, William Pence, David Valls-Gabaud, Christopher Willmer, for their useful and constructive comments and suggestions. Finally we should thank all the (sometimes anonymous) developers in various online forums which patiently answered all our small (but important) technical questions. All work on Gnuastro has been voluntary, but we are most grateful to the following institutions (in chronological order) for hosting us in our research:

Ministry of education, culture, sports, science and technology (MEXT), Japan.
Tohoku University Astronomical Institute, Sendai, Japan.
University of Salento, Lecce, Italy.
Centre national de la recherche scientifique (CNRS), France.
Centre de Recherche Astrophysique de Lyon, University of Lyon 1, France.

2 Tutorials

In this chapter we give several tutorials or cookbooks on how to use the various tools in Gnuastro for your scientific purposes. In these tutorials, we have intentionally avoided too many cross references to make it more easily readable. To get more information about a particular program, you can visit the section with the same name as the program in this book. Each program section starts by explaining the general concepts behind what it does. If you only want to see an explanation of the options and arguments of any program, see the subsection titled 'Invoking ProgramName'. See Section 1.10 [Conventions], page 12, for an explanation of the conventions we use in the example codes through the book.

The tutorials in this section use a fictional setting of some historical figures in the history of astronomy. We have tried to show how Gnuastro would have been helpful for them in making their discoveries if there were GNU/Linux computers in their times! Please excuse us for any historical inaccuracy, this is not intended to be a historical reference. This form of presentation can make the tutorials more pleasant and entertaining to read while also being more practical (explaining from a user's point of view)[1]. The main reference for the historical facts mentioned in these fictional settings was Wikipedia.

2.1 Hubble visually checks and classifies his catalog

In 1924 Hubble[2] announced his discovery that some of the known nebulous objects are too distant to be within the the Milky Way (or Galaxy) and that they were probably distant Galaxies[3] in their own right. He had also used them to show that the redshift of the nebulae increases with their distance. So now he wants to study them more accurately to see what they actually are. Since they are nebulous or amorphous, they can't be modeled (like stars that are always a point) easily. So there is no better way to distinguish them than to visually inspect them and see if it is possible to classify these nebulae or not.

Hubble has stored all the FITS images of the objects he wants to visually inspect in his `/mnt/data/images` directory. He has also stored his catalog of extra-galactic nebulae in `/mnt/data/catalogs/extragalactic.txt`. Any normal user on his GNU/Linux system (including himself) only has read access to the contents of the `/mnt/data` directory. He has done this by running this command as root:

```
# chmod -R 755 /mnt/data
```

[1] This form of presenting a tutorial was influenced by the PGF/TikZ and Beamer manuals. The first provides graphic capabilities, while with the second you can make presentation slides in TEX and LATEX. In these manuals, Till Tantau (author of the manual) uses Euclid as the protagonist. There are also some nice words of wisdom for Unix-like systems called "Rootless Root": `http://catb.org/esr/writings/unix-koans/`. These also have a similar style but they use a mythical figure named Master Foo. If you already have some experience in Unix-like systems, you will definitely find these "Unix Koans" very entertaining.

[2] Edwin Powell Hubble (1889 – 1953 A.D.) was an American astronomer who can be considered as the father of extra-galactic astronomy, by proving that some nebulae are too distant to be within the Galaxy. He then went on to show that the universe appears to expand and also done a visual classification of the galaxies that is known as the Hubble fork.

[3] Note that at that time, "Galaxy" was a proper noun used to refer to the Milky way. The concept of a galaxy as we define it today had not yet become common. Hubble played a major role in creating today's concept of a galaxy.

Hubble has done this intentionally to avoid mistakenly deleting or modifying the valuable images he has taken at Mount Wilson while he is working as an ordinary user. Retaking all those images and data is simply not an option. In fact they are also in another hard disk (`/dev/sdb1`). So if the hard disk which stores his GNU/Linux distribution suddenly malfunctions due to work load, his data is not in harms way. That hard disk is only mounted to this directory when he wants to use it with the command:

```
# mount /dev/sdb1 /mnt/data
```

In short, Hubble wants to keep his data safe and fortunately by default Gnuastro allows for this. Hubble creates a temporary `visualcheck` directory in his home directory for this check. He runs the following commands to make the directory and change to it[4]:

```
$ mkdir ~/visualcheck
$ cd ~/visualcheck
$ pwd
/home/edwin/visualcheck
$ ls
```

Hubble has multiple images in `/mnt/data/images`, some of his targets might be on the edges of an image and so several images need to be stitched to give a good view of them. Also his extra-galactic targets belong to various pointings in the sky, so they are not in one large image. Gnuastro's Crop is just the program he wants. The catalog in `extragalactic.txt` is a plain text file which stores the basic information of all his known 200 extra-galactic nebulae. In its second column it has each object's Right Ascension (the first column is a label he has given to each object) and in the third the object's declination.

```
$ astcrop --racol=2 --deccol=3 /mnt/data/images/*.fits     \
             /mnt/data/catalogs/extragalactic.txt
Crop started on Tue Jun  14 10:18:11 1932
  ---- ./4_crop.fits                  1 1
  ---- ./2_crop.fits                  1 1
  ---- ./1_crop.fits                  1 1
[[[ Truncated middle of list ]]]
  ---- ./198_crop.fits                1 1
  ---- ./195_crop.fits                1 1
  - 200 images created.
  - 200 were filled in the center.
  - 0 used more than one input.
Crop finished in:  2.429401 (seconds)
```

Hubble already knows that thread allocation to the the CPU cores is asynchronous. Hence each time you run it, the order of which job gets done first differs. When using Crop the order of outputs is irrelevant since each crop is independent of the rest. This is why the crops are not necessarily created in the same input order. He is satisfied with the default width of the outputs (which he inspected by running `$ astcrop -P`). If he wanted a different width for the cropped images, he could do that with the `--wwidth` option which accepts a value in arc-seconds. When he lists the contents of the directory again he finds his 200 objects as separate FITS images.

[4] The `pwd` command is short for "Print Working Directory" and `ls` is short for "list" which shows the contents of a directory.

```
$ ls
1_crop.fits 2_crop.fits ... 200_crop.fits
```

The FITS image format was not designed for efficient/fast viewing, but mainly for accurate storing of the data. So he chooses to convert the cropped images to a more common image format to view them more quickly and easily through standard image viewers (which load much faster than FITS image viewer). JPEG is one of the most recognized image formats that is supported by most image viewers. Fortunately Gnuastro has just such a tool to convert various types of file types to and from each other: ConvertType. Hubble has already heard of GNU Parallel from one of his colleagues at Mount Wilson Observatory. It allows multiple instances of a command to be run simultaneously on the system, so he uses it in conjunction with ConvertType to convert all the images to JPEG.

```
$ parallel astconvertt -ojpg ::: *_crop.fits
```

For his graphical user interface Hubble is using GNOME which is the default in most distributions in GNU/Linux. The basic image viewer in GNOME is the Eye of GNOME, which has the executable file name `eog`[5]. Since he has used it before, he knows that once it opens an image, he can use the `ENTER` or `SPACE` keys on the keyboard to go to the next image in the directory or the `Backspace` key to to go the previous image. So he opens the image of the first object with the command below and with his cup of coffee in his other hand, he flips through his targets very fast to get a good initial impression of the morphologies of these extra-galactic nebulae.

```
$ eog 1_crop.jpg
```

Hubble's cup of coffee is now finished and he also got a nice general impression of the shapes of the nebulae. He tentatively/mentally classified the objects into three classes while doing the visual inspection. One group of the nebulae have a very simple elliptical shape and seem to have no internal special structure, so he gives them code 1. Another clearly different class are those which have spiral arms which he associates with code 2 and finally there seems to be a class of nebulae in between which appear to have a disk but no spiral arms, he gives them code 3.

Now he wants to know how many of the nebulae in his extra-galactic sample are within each class. Repeating the same process above and writing the results on paper is very time consuming and prone to errors. Fortunately Hubble knows the basics of GNU Bash shell programming, so he writes the following short script with a loop to help him with the job. After all, computers are made for us to operate and knowing basic shell programming gives Hubble this ability to creatively operate the computer as he wants. So using GNU Emacs[6] (his favorite text editor) he puts the following text in a file named `classify.sh`.

```
for name in *.jpg
do
    eog $name &
    processid=$!
    echo -n "$name belongs to class: "
    read class
```

[5] Eye of GNOME is only available for users of the GNOME graphical desktop environment which is the default in most GNU/Linux distributions. If you use another graphical desktop environment, replace `eog` with any other image viewer.

[6] This can be done with any text editor

```
        echo $name $class >> classified.txt
        kill $processid
    done
```

Fortunately GNU Emacs or even simpler editors like Gedit (part of the GNOME graphical user interface) will display the variables and shell constructs in different colors which can really help in understanding the script. Put simply, the `for` loop gets the name of each JPEG file in the directory this script is run in and puts it in `name`. In the shell, the value of a variable is used by putting a `$` sign before the variable name. Then Eye of GNOME is run on the image in the background to show him that image and its process ID is saved internally (this is necessary to close Eye of GNOME later). The shell then prompts the user to specify a class and after saving it in `class`, it prints the file name and the given class in the next line of a file named `classified.txt`. To make the script executable (so he can run it later any time he wants) he runs:

```
$ chmod +x classify.sh
```

Now he is ready to do the classification, so he runs the script:

```
$ ./classify.sh
```

In the end he can delete all the JPEG and FITS files along with Crop's log file with the following short command. The only files remaining are the script and the result of the classification.

```
$ rm *.jpg *.fits astcrop.txt
$ ls
classified.txt   classify.sh
```

He can now use `classified.txt` as input to a plotting program to plot the histogram of the classes and start making interpretations about what these nebulous objects that are outside of the Galaxy are.

2.2 Sufi simulates a detection

It is the year 953 A.D. and Sufi[7] is in Shiraz as a guest astronomer. He had come there to use the advanced 123 centimeter astrolabe for his studies on the Ecliptic. However, something was bothering him for a long time. While mapping the constellations, there were several non-stellar objects that he had detected in the sky, one of them was in the Andromeda constellation. During a trip he had to Yemen, Sufi had seen another such object in the southern skies looking over the Indian ocean. He wasn't sure if such cloud-like non-stellar objects (which he was the first to call 'Sahābi' in Arabic or 'nebulous') were real astronomical objects or if they were only the result of some bias in his observations. Could such diffuse objects actually be detected at all with his detection technique?

He still had a few hours left until nightfall (when he would continue his studies on the ecliptic) so he decided to find an answer to this question. He had thoroughly studied Claudius Ptolemy's (90 – 168 A.D) Almagest and had made lots of corrections to it, in particular in measuring the brightness. Using his same experience, he was able to measure a magnitude for the objects and wanted to simulate his observation to see if a simulated

[7] Abd al-rahman Sufi (903 – 986 A.D.), also known in Latin as Azophi was an Iranian astronomer. His manuscript "Book of fixed stars" contains the first recorded observations of the Andromeda galaxy, the Large Magellanic Cloud and seven other non-stellar or 'nebulous' objects.

object with the same brightness and size could be detected in a simulated noise with the same detection technique. The general outline of the steps he wants to take are:

1. Make some mock profiles in an oversampled image. The initial mock image has to be oversampled prior to convolution or other forms of transformation in the image. Through his experiences, Sufi knew that this is because the image of heavenly bodies is actually transformed by the atmosphere or other sources outside the atmosphere (for example gravitational lenses) prior to being sampled on an image. Since that transformation occurs on a continuous grid, to best approximate it, he should do all the work on a finer pixel grid. In the end he can re-sample the result to the initially desired grid size.
2. Convolve the image with a PSF image that is oversampled to the same value as the mock image. Since he wants to finish in a reasonable time and the PSF kernel will be very large due to oversampling, he has to use frequency domain convolution which has the side effect of dimming the edges of the image. So in the first step above he also has to build the image to be larger by at least half the width of the PSF convolution kernel on each edge.
3. With all the transformations complete, the image should be re-sampled to the same size of the pixels in his detector.
4. He should remove those extra pixels on all edges to remove frequency domain convolution artifacts in the final product.
5. He should add noise to the (until now, noise-less) mock image. After all, all observations have noise associated with them.

Fortunately Sufi had heard of GNU Astronomy Utilities from a colleague in Isfahan (where he worked) and had installed it on his computer a year before. It had tools to do all the steps above. He had used MakeProfiles before, but wasn't sure which columns he had chosen in his user or system wide configuration files for which parameters, see Section 4.2 [Configuration files], page 59. So to start his simulation, Sufi runs MakeProfiles with the `-P` option to make sure what columns in a catalog MakeProfiles currently recognizes and the output image parameters. In particular, Sufi is interested in the recognized columns (shown below).

```
$ astmkprof -P

[[[ ... Truncated lines ... ]]]

# Output:
 type         float32     # Type of output: e.g., int16, float32, etc...
 naxis1       1000        # Number of pixels along first FITS axis.
 naxis2       1000        # Number of pixels along second FITS axis.
 oversample   5           # Scale of oversampling (>0 and odd).

[[[ ... Truncated lines ... ]]]

# Columns, by info (see ‘--searchin’), or number (starting from 1):
 xcol         2           # Center along first FITS axis (horizontal).
 ycol         3           # Center along second FITS axis (vertical).
```

```
fcol        4           # sersic (1), moffat (2), gaussian (3),
                        # point (4), flat (5), circumference (6).
rcol        5           # Effective radius or FWHM in pixels.
ncol        6           # Sersic index or Moffat beta.
pcol        7           # Position angle.
qcol        8           # Axis ratio.
mcol        9           # Magnitude.
tcol        10          # Truncation in units of radius or pixels.

[[[ ... Truncated lines ... ]]]
```

In Gnuastro, column counting starts from 1, so the columns are ordered such that the first column (number 1) can be an ID he specifies for each object (and MakeProfiles ignores), each subsequent column is used used for another property of the profile. It is also possible to use column names for the values of these options and change these defaults, but Sufi preferred to stick to the defaults. Fortunately MakeProfiles has the capability to also make the PSF which is to be used on the mock image and using the `--prepforconv` option, he can also make the mock image to be larger by the correct amount and all the sources to be shifted by the correct amount.

For his initial check he decides to simulate the nebula in the Andromeda constellation. The night he was observing, the PSF had roughly a FWHM of about 5 pixels, so as the first row (profile), he defines the PSF parameters and sets the radius column (`rcol` above, fifth column) to `5.000`, he also chooses a Moffat function for its functional form. Remembering how diffuse the nebula in the Andromeda constellation was, he decides to simulate it with a mock Sérsic index 1.0 profile. He wants the output to be 500 pixels by 500 pixels, so he puts the mock profile in the center. Looking at his drawings of it, he decides a reasonable effective radius for it would be 40 pixels on this image pixel scale, he sets the axis ratio and position angle to approximately correct values too and finally he sets the total magnitude of the profile to 3.44 which he had accurately measured. Sufi also decides to truncate both the mock profile and PSF at 5 times the respective radius parameters. In the end he decides to put four stars on the four corners of the image at very low magnitudes as a visual scale.

Using all the information above, he creates the catalog of mock profiles he wants in a file named `cat.txt` (short for catalog) using his favorite text editor and stores it in a directory named `simulationtest` in his home directory. [The `cat` command prints the contents of a file, short for concatenation. So please copy-paste the lines after "`cat cat.txt`" into `cat.txt` when the editor opens in the steps above it, note that there are 7 lines, first one starting with `#`]:

```
$ mkdir ~/simulationtest
$ cd ~/simulationtest
$ pwd
/home/rahman/simulationtest
$ emacs cat.txt
$ ls
cat.txt
$ cat cat.txt
# Column 4: PROFILE_NAME [,str7] Radial profile's functional name
```

```
1  0.0000   0.0000  moffat  5.000  4.765  0.0000  1.000  30.000  5.000
2  250.00   250.00  sersic  40.00  1.000  -25.00  0.400  3.4400  5.000
3  50.000   50.000  point   0.000  0.000  0.0000  0.000  9.0000  0.000
4  450.00   50.000  point   0.000  0.000  0.0000  0.000  9.2500  0.000
5  50.000   450.00  point   0.000  0.000  0.0000  0.000  9.5000  0.000
6  450.00   450.00  point   0.000  0.000  0.0000  0.000  9.7500  0.000
```

The zero-point magnitude for his observation was 18. Now he has all the necessary parameters and runs MakeProfiles with the following command:

```
$ astmkprof --prepforconv --naxis1=500 --naxis2=500         \
            --zeropoint=18.0 cat.txt
MakeProfiles started on Sat Oct  6 16:26:56 953
  - 6 profiles read from cat.txt
  - Random number generator (RNG) type: mt19937
  - Using 8 threads.
  ---- row 2 complete, 5 left to go
  ---- row 3 complete, 4 left to go
  ---- row 4 complete, 3 left to go
  ---- row 5 complete, 2 left to go
  ---- ./0_cat.fits created.
  ---- row 0 complete, 1 left to go
  ---- row 1 complete, 0 left to go
  - ./cat.fits created.                                0.041651 seconds
MakeProfiles finished in 0.267234 seconds

$ls
0_cat.fits  cat.fits  cat.txt
```

The file `0_cat.fits` is the PSF Sufi had asked for and `cat.fits` is the image containing the other 5 objects. The PSF is now available to him as a separate file for the convolution step. While he was preparing the catalog, one of his students approached him and was also following the steps. When he opened the image, the student was surprised to see that all the stars are only one pixel and not in the shape of the PSF as we see when we image the sky at night. So Sufi explained to him that the stars will take the shape of the PSF after convolution and this is how they would look if we didn't have an atmosphere or an aperture when we took the image. The size of the image was also surprising for the student, instead of 500 by 500, it was 2630 by 2630 pixels. So Sufi had to explain why oversampling is very important for parts of the image where the flux change is significant over a pixel. Sufi then explained to him that after convolving we will re-sample the image to get our originally desired size. To convolve the image, Sufi ran the following command:

```
$ astconvolve --kernel=0_cat.fits cat.fits
Convolve started on Mon Apr  6 16:35:32 953
  - Using 8 CPU threads.
  - Input: cat.fits (hdu: 1)
  - Kernel: 0_cat.fits (hdu: 1)
  - Input and Kernel images padded.                    0.075541 seconds
  - Images converted to frequency domain.              6.728407 seconds
```

```
  - Multiplied in the frequency domain.               0.040659 seconds
  - Converted back to the spatial domain.             3.465344 seconds
  - Padded parts removed.                             0.016767 seconds
Convolve finished in:  10.422161 seconds

$ls
0_cat.fits  cat_convolved.fits  cat.fits  cat.txt
```

When convolution finished, Sufi opened the `cat_convolved.fits` file and showed the effect of convolution to his student and explained to him how a PSF with a larger FWHM would make the points even wider. With the convolved image ready, they were prepared to re-sample it to the original pixel scale Sufi had planned [from the `$ astmkprof -P` command above, recall that MakeProfiles had oversampled the image by 5 times]. Sufi explained the basic concepts of warping the image to his student and ran Warp with the following command:

```
$ astwarp --scale=1/5 --centeroncorner cat_convolved.fits
Warp started on Mon Apr  6 16:51:59 953
 Using 8 CPU threads.
 Input: cat_convolved.fits (hdu: 1)
 matrix:
0.2000   0.0000   0.4000
0.0000   0.2000   0.4000
0.0000   0.0000   1.0000

$ ls
0_cat.fits          cat_convolved_scaled.fits     cat.txt
cat_convolved.fits  cat.fits

$ astfits -p cat_convolved_scaled.fits | grep NAXIS
NAXIS   =                    2 / number of data axes
NAXIS1  =                  526 / length of data axis 1
NAXIS2  =                  526 / length of data axis 2
```

`cat_convolved_warped.fits` now has the correct pixel scale. However, the image is still larger than what we had wanted, it is 526 ($500 + 13 + 13$) by 526 pixels. The student is slightly confused, so Sufi also resamples the PSF with the same scale and shows him that it is 27 ($2 \times 13 + 1$) by 27 pixels. Sufi goes on to explain how frequency space convolution will dim the edges and that is why he added the `--prepforconv` option to MakeProfiles, see Section 8.1.2 [If convolving afterwards], page 200. Now that convolution is done, Sufi can remove those extra pixels using Crop with the command below. Crop's `--section` option accepts coordinates inclusively and counting from 1 (according to the FITS standard), so the crop's first pixel has to be 14, not 13.

```
$ astcrop cat_convolved_scaled.fits --section=14:*-13,14:*-13    \
          --zeroisnotblank
Crop started on Sat Oct  6 17:03:24 953
  - Read metadata of 1 image.                         0.001304 seconds
  ---- ...nvolved_scaled_cropped.fits created: 1 input.
Crop finished in:  0.027204 seconds
```

```
$ls
0_cat.fits          cat_convolved_scaled_cropped.fits  cat.fits
cat_convolved.fits  cat_convolved_scaled.fits          cat.txt
```

Finally, `cat_convolved_scaled_cropped.fits` has the same dimensions as Sufi had desired in the beginning. All this trouble was certainly worth it because now there is no dimming on the edges of the image and the profile centers are more accurately sampled. The final step to simulate a real observation would be to add noise to the image. Sufi set the zeropoint magnitude to the same value that he set when making the mock profiles and looking again at his observation log, he had measured the background flux near the nebula had a magnitude of 7 that night. So using these values he ran MakeNoise:

```
$ astmknoise --zeropoint=18 --background=7 --output=out.fits    \
             cat_convolved_warped_crop.fits
MakeNoise started on Mon Apr  6 17:05:06 953
  - Generator type: mt19937
  - Generator seed: 1428318100
MakeNoise finished in:  0.033491 (seconds)

$ls
0_cat.fits         cat_convolved_scaled_cropped.fits cat.fits  out.fits
cat_convolved.fits cat_convolved_scaled.fits         cat.txt
```

The `out.fits` file now contains the noised image of the mock catalog Sufi had asked for. Seeing how the `--output` option allows the user to specify the name of the output file, the student was confused and wanted to know why Sufi hadn't used it before? Sufi then explained to him that for intermediate steps it is best to rely on the automatic output, see Section 4.8 [Automatic output], page 77. Doing so will give all the intermediate files the same basic name structure, so in the end you can simply remove them all with the Shell's capabilities. So Sufi decided to show this to the student by making a shell script from the commands he had used before.

The command-line shell has the capability to read all the separate input commands from a file. This is very useful when you want to do the same thing multiple times, with only the names of the files or minor parameters changing between the different instances. Using the shell's history (by pressing the up keyboard key) Sufi reviewed all the commands and then he retrieved the last 5 commands with the `$ history 5` command. He selected all those lines he had input and put them in a text file named `mymock.sh`. Then he defined the `edge` and `base` shell variables for easier customization later. Finally, before every command, he added some comments (lines starting with `#`) for future readability.

```
# Basic settings:
edge=13
base=cat

# Remove any existing image to avoid confusion.
rm out.fits

# Run MakeProfiles to create an oversampled FITS image.
astmkprof --prepforconv --naxis1=500 --naxis2=500               \
```

```
          --zeropoint=18.0 "$base".txt

# Convolve the created image with the kernel.
astconvolve --kernel=0_"$base".fits "$base".fits

# Scale the image back to the intended resolution.
astwarp --scale=1/5 --centeroncorner "$base"_convolved.fits

# Crop the edges out (dimmed during convolution). '--section' accepts
# inclusive coordinates, so the start of start of the section must be
# one pixel larger than its end.
st_edge=$(( edge + 1 ))
astcrop "$base"_convolved_scaled.fits --zeroisnotblank          \
        --section=$st_edge:*-$edge,$st_edge:*-$edge

# Add noise to the image.
astmknoise --zeropoint=18 --background=7 --output=out.fits      \
           "$base"_convolved_scaled_cropped.fits

# Remove all the temporary files.
rm 0*.fits cat*.fits
```

He used this chance to remind the student of the importance of comments in code or shell scripts: when writing the code, you have a very good mental picture of what you are doing, so writing comments might seem superfluous and excessive. However, in one month when you want to re-use the script, you have lost that mental picture and rebuilding it is can be very time-consuming and frustrating. The importance of comments is further amplified when you want to share the script with a friend/colleague. So it is very good to accompany any script/code with useful comments while you are writing it (have a good mental picture of what/why you are doing something).

Sufi then explained to the eager student that you define a variable by giving it a name, followed by an `=` sign and the value you want. Then you can reference that variable from anywhere in the script by calling its name with a `$` prefix. So in the script whenever you see `$base`, the value we defined for it above is used. If you use advanced editors like GNU Emacs or even simpler ones like Gedit (part of the GNOME graphical user interface) the variables will become a different color which can really help in understanding the script. We have put all the `$base` variables in double quotation marks (`"`) so the variable name and the following text do not get mixed, the shell is going to ignore the `"` after replacing the variable value. To make the script executable, Sufi ran the following command:

```
$ chmod +x mymock.sh
```

Then finally, Sufi ran the script, simply by calling its file name:

```
$ ./mymock.sh
```

After the script finished, the only file remaining is the `out.fits` file that Sufi had wanted in the beginning. Sufi then explained to the student how he could run this script anywhere that he has a catalog if the script is in the same directory. The only thing the student had to modify in the script was the name of the catalog (the value of the `base` variable in the

start of the script) and the value to the `edge` variable if he changed the PSF size. The student was also very happy to hear that he won't need to make it executable again when he makes changes later, it will remain executable unless he explicitly changes the executable flag with `chmod`.

The student was really excited, since now, through simple shell scripting, he could really speed up his work and run any command in any fashion he likes allowing him to be much more creative in his works. Until now he was using the graphical user interface which doesn't have such a facility and doing repetitive things on it was really frustrating and some times he would make mistakes. So he left to go and try scripting on his own computer.

Sufi could now get back to his own work and see if the simulated nebula which resembled the one in the Andromeda constellation could be detected or not. Although it was extremely faint[8], fortunately it passed his detection tests and he wrote it in the draft manuscript that would later become "Book of fixed stars". He still had to check the other nebula he saw from Yemen and several other such objects, but they could wait until tomorrow (thanks to the shell script, he only has to define a new catalog). It was nearly sunset and they had to begin preparing for the night's measurements on the ecliptic.

[8] The brightness of a diffuse object is added over all its pixels to give its final magnitude, see Section 8.1.3 [Flux Brightness and magnitude], page 201. So although the magnitude 3.44 (of the mock nebula) is orders of magnitude brighter than 6 (of the stars), the central galaxy is much fainter. Put another way, the brightness is distributed over a large area in the case of a nebula.

3 Installation

The latest released version of Gnuastro source code is always available at the following URL:

```
http://ftpmirror.gnu.org/gnuastro/gnuastro-latest.tar.gz
```

Section 1.1 [Quick start], page 1, discribes the commands necessary to configure, build, and install Gnuastro on your system. This chapter will be useful in cases where the simple procedure above is not sufficient, for example your system lacks a mandatory/optional dependency (in other words, you can't pass the `$ ./configure` step), or you want greater customization, or you want to build and install Gnuastro from other random points in its history, or you want a higher level of control on the installation. Thus if you were happy with downloading the tarball and following Section 1.1 [Quick start], page 1, then you can safely ignore this chapter and come back to it in the future if you need more customization.

Section 3.1 [Dependencies], page 25, describes the mandatory, optional and bootstrapping dependencies of Gnuastro. Only the first group are required/mandatory when you are building Gnuastro using a tarball (see Section 3.2.1 [Release tarball], page 31), they are very basic and low-level tools used in most astronomical software, so you might already have them installed, if not they are very easy to install as described for each. Section 3.2 [Downloading the source], page 31, discusses the two methods you can obtain the source code: as a tarball (a significant snapshot in Gnuastro's history), or the full history[1]. The latter allows you to build Gnuastro at any random point in its history (for example to get bug fixes or new features that are not released as a tarball yet).

The building and installation of Gnuastro is heavily customizable, to learn more about them, see Section 3.3 [Build and install], page 36. This section is essentially a thorough explanation of the steps in Section 1.1 [Quick start], page 1. It discusses ways you can influence the building and installation. If you encounter any problems in the installation process, it is probably already explained in Section 3.3.4 [Known issues], page 46. In Appendix B [Other useful software], page 343, the installation and usage of some other free software that are not directly required by Gnuastro but might be useful in conjunction with it is discussed.

3.1 Dependencies

The dependencies needed to build and install Gnuastro are defined by the features you want and how you would like to obtain the source code (see Section 3.2 [Downloading the source], page 31). A minimal set of dependencies are mandatory, if they are not present you cannot get passed the configuration step. Such mandatory dependencies are therefore very basic (low-level) tools which are easy to obtain, build and install, see Section 3.1.1 [Mandatory dependencies], page 26, for a full discussion.

If you have the packages of Section 3.1.2 [Optional dependencies], page 28, Gnuastro will have additional functionality (for example converting FITS images to JPEG or PDF). If you are installing from a tarball as explained in Section 1.1 [Quick start], page 1, you can stop reading after this section. However, if you decided to use the version controlled source instead of the tarball (see Section 3.2.2 [Version controlled source], page 32), an additional bootstrapping step is required before configuration and its dependencies are explained in Section 3.1.3 [Bootstrapping dependencies], page 29.

---

[1] Section 3.1.3 [Bootstrapping dependencies], page 29, are required if you clone the full history.

3.1.1 Mandatory dependencies

The mandatory Gnuastro dependencies are very basic and low-level tools. They all follow the same basic GNU based build system (like that shown in Section 1.1 [Quick start], page 1), so even if you don't have them, installing them should be pretty straightforward. In this section we explain each program and any specific note that might be necessary in the installation.

The most basic choice is to build the packages from source yourself, instead of relying on your distribution's package management system. While the latter choice is indeed possible, we recommend that you build these dependencies yourself as discussed below. We will send out notifications in the `info-gnuastro` mailing list, see Section 1.9 [Announcements], page 11, when we find out that these requirements are updated.

1. For each package, Gnuastro might preform better (or require) certain configuration options that your distribution's package managers didn't add for you. If present, these configuration options are explained during the installation of each in the sections below. When the proper configuration has not been set, the programs should complain and inform you.
2. Your distribution's pre-built package might not be the most recent release.
3. For the libraries, they might separate the binary file from the header files, see Section 3.3.4 [Known issues], page 46.
4. Like any other tool, the science you derive from Gnuastro's tools highly depend on these lower level dependencies, so generally it is much better to have a close connection with them. By reading their manuals, installing them and staying up to date with changes/bugs in them, your scientific results and understanding will also correspondingly improve.

3.1.1.1 GNU Scientific library

The GNU Scientific Library (`http://www.gnu.org/software/gsl/`), or GSL, is a large collection of functions that are very useful in scientific applications, for example integration, random number generation, and Fast Fourier Transform among many others. To install GSL from source, you can run the following commands after you have downloaded `gsl-latest.tar.gz` (`ftp://ftp.gnu.org/gnu/gsl/gsl-latest.tar.gz`):

```
$ tar xf gsl-latest.tar.gz
$ cd gsl-X.X                      # Replace X.X with version number.
$ ./configure
$ make -j8                        # Replace 8 with no. CPU threads.
$ make check
$ sudo make install
```

3.1.1.2 CFITSIO

CFITSIO (`http://heasarc.gsfc.nasa.gov/fitsio/`) is the closest you can get to the pixels in a FITS image while remaining faithful to the FITS standard (`http://fits.gsfc.nasa.gov/fits_standard.html`). It is written by William Pence, the principal author of

the FITS standard[2], and is regularly updated. Setting the definitions for all other software packages using FITS images.

Some GNU/Linux distributions have CFITSIO in their package managers, if it is available and updated, you can use it. One problem that might occur is that CFITSIO might not be configured with the `--enable-reentrant` option by the distribution. This option allows CFITSIO to open a file in multiple threads, it can thus provide great speed improvements. If CFITSIO was not configured with this option, any program which needs this capability will warn you and abort when you ask for multiple threads (see Section 4.3 [Multi-threaded operations], page 62).

To install CFITSIO from source, we strongly recommend that you have a look through Chapter 2 (Creating the CFITSIO library) of the CFITSIO manual and understand the options you can pass to `$ ./configure` (they aren't too much). This is a very basic package for most astronomical software and it is best that you configure it nicely with your system. Once you download the source and unpack it, the following configure script should be enough for most purposes. Don't forget to read chapter two of the manual though, for example the second option is only for 64bit systems. The manual also explains how to check if it has been installed correctly.

CFITSIO comes with two executables called fpack and funpack. From their manual: they "are standalone programs for compressing and uncompressing images and tables that are stored in the FITS (Flexible Image Transport System) data format. They are analogous to the gzip and gunzip compression programs except that they are optimized for the types of astronomical images that are often stored in FITS format". The commands below will compile and install them on your system along with CFITSIO. They are not essential for Gnuastro, since they are just wrappers for functions within CFITSIO, but they can come in handy. The `make utils` command is only available for versions above 3.39, it will build these executables along with several other test executables which are deleted before the installation (otherwise they will also be installed).

The CFITSIO installation from source process is given below. Let's assume you have downloaded `cfitsio_latest.tar.gz` (`http://heasarc.gsfc.nasa.gov/FTP/software/fitsio/c/cfitsio_latest.tar.gz`) and are in the same directory:

```
$ tar xf cfitsio_latest.tar.gz
$ cd cfitsio
$ ./configure --prefix=/usr/local --enable-sse2 --enable-reentrant
$ make -j8                          # Replace 8 with no. CPU threads.
$ make utils
$ ./testprog > testprog.lis
$ diff testprog.lis testprog.out    # Should have no output
$ cmp testprog.fit testprog.std     # Should have no output
$ rm cookbook fitscopy imcopy smem speed testprog
$ sudo make install
```

[2] Pence, W.D. et al. Definition of the Flexible Image Transport System (FITS), version 3.0. (2010) Astronomy and Astrophysics, Volume 524, id.A42, 40 pp.

3.1.1.3 WCSLIB

WCSLIB (http://www.atnf.csiro.au/people/mcalabre/WCS/) is written and maintained by one of the authors of the World Coordinate System (WCS) definition in the FITS standard (http://fits.gsfc.nasa.gov/fits_standard.html)[3], Mark Calabretta. It might be already built and ready in your distribution's package management system. However, here the installation from source is explained, for the advantages of installation from source please see Section 3.1.1 [Mandatory dependencies], page 26. To install WCSLIB you will need to have CFITSIO already installed, see Section 3.1.1.2 [CFITSIO], page 26.

WCSLIB also has plotting capabilities which use PGPLOT (a plotting library for C). If you wan to use those capabilities in WCSLIB, Section B.2 [PGPLOT], page 344, provides the PGPLOT installation instructions. However PGPLOT is old[4], so its installation is not easy, there are also many great modern WCS plotting tools (mostly in written in Python). Hence, if you will not be using those plotting functions in WCSLIB, you can configure it with the `--without-pgplot` option as shown below. Let's assume you have downloaded `wcslib.tar.bz2` (ftp://ftp.atnf.csiro.au/pub/software/wcslib/wcslib.tar.bz2) and are in the same directory:

```
$ tar xf wcslib.tar.bz2
$ cd wcslib-X.X                       # Replace X.X with version number
$ ./configure --without-pgplot LIBS="-pthread -lm"
$ make -j8                            # Replace 8 with no. CPU threads.
$ make check
$ sudo make install
```

3.1.2 Optional dependencies

The libraries listed here are only used for very specific applications, therefore if you don't want these operations, Gnuastro will be built and installed without them and you don't have to have the dependencies.

If the `./configure` script can't find these requirements, it will warn you in the end that they are not present and notify you of the operation(s) you can't do due to not having them. If the output you request from a program requires a missing library, that program is going to warn you and abort. In the case of executables like GPL GhostScript, if you install them at a later time, the program will run. This is because if required libraries are not present at build time, the executables cannot be built, but an executable is called by the built program at run time so if it becomes available, it will be used. If you do install an optional library later, you will have to rebuild Gnuastro and reinstall it for it to take effect.

libgit2
: Git is one of the most common version control systems (see Section 3.2.2 [Version controlled source], page 32). When `libgit2` is present, and Gnuastro's programs are run within a version controlled directory, outputs will contain the version number of the working directory's repository for future reproducibility. See the `COMMIT` keyword header in Section 4.9 [Output headers], page 78, for a discussion.

---

[3] Greisen E.W., Calabretta M.R. (2002) Representation of world coordinates in FITS. Astronomy and Astrophysics, 395, 1061-1075.

[4] As of early June 2016, its most recent version was uploaded in February 2001.

libjpeg
: libjpeg is only used by ConvertType to read from and write to JPEG images. libjpeg (`http://www.ijg.org/`) is a very basic library that provides tools to read and write JPEG images, most of the GNU/Linux graphic programs and libraries use it. Therefore you most probably already have it installed. libjpeg-turbo (`http://libjpeg-turbo.virtualgl.org/`) is an alternative to libjpeg. It uses SIMD instructions for ARM based systems that significantly decreases the processing time of JPEG compression and decompression algorithms.

GPL Ghostscript
: GPL Ghostscript's executable (`gs`) is called used by ConvertType to compile a PDF file from a source PostScript file, see Section 5.2 [ConvertType], page 87. Therefore its headers (and libraries) are not needed. With a very high probability you already have it in your GNU/Linux distribution. Unfortunately it does not follow the standard GNU build style so installing it is very hard. It is best to rely on your distribution's package managers for this.

3.1.3 Bootstrapping dependencies

Bootstrapping is only necessary if you have decided to obtain the full version controlled history of Gnuastro, see Section 3.2.2 [Version controlled source], page 32, and Section 3.2.2.1 [Bootstrapping], page 33. Using the version controlled source enables you to always be up to date with the most recent development work of Gnuastro (bug fixes, new functionalities, improved algorithms and etc). If you have downloaded a tarball (see Section 3.2 [Downloading the source], page 31), then you can ignore this subsection.

To successfully run the bootstrapping process, there are some additional dependencies to those discussed in the previous subsections. These are low level tools that are used by a large collection of Unix-like operating systems programs, therefore they are most probably already available in your system. If they are not already installed, you should be able to easily find them in any GNU/Linux distribution package management system (`apt-get`, `yum`, `pacman` and etc). The short names in parenthesis in `typewriter` font after the package name can be used to search for them in your package manager. For the GNU Portability Library, GNU Autoconf Archive and TEX Live, it is recommended to use the instructions here, not your operating system's package manager.

GNU Portability Library (Gnulib)
: To ensure portability for a wider range of operating systems (those that don't include GNU C library, namely glibc), Gnuastro depends on the GNU portability library, or Gnulib. Gnulib keeps a copy of all the functions in glibc, implemented (as much as possible) to be portable to other operating systems. The `bootstrap` script can automatically clone Gnulib (as a `gnulib/` directory inside Gnuastro), however, as described in Section 3.2.2.1 [Bootstrapping], page 33, this is not recommended.

: The recommended way to bootstrap Gnuastro is to first clone Gnulib and the Autoconf archives (see below) into a local directory outside of Gnuastro. Let's

call it `DEVDIR`[5] (which you can set to any directory). Currently in Gnuastro, both Gnulib and Autoconf archives have to be cloned in the same top directory[6] like the case here[7]:

```
$ DEVDIR=/home/yourname/Development
$ cd $DEVDIR
$ git clone git://git.sv.gnu.org/gnulib.git
$ git clone git://git.sv.gnu.org/autoconf-archive.git
```

You now have the full version controlled source of these two repositories in separate directories. Both these packages are regularly updated, so every once in a while, you can run `$ git pull` within them to get any possible updates.

GNU Automake (`automake`)
: GNU Automake will build the `Makefile.in` files in each sub-directory using the (hand-written) `Makefile.am` files. The `Makefile.in`s are subsequently used to generate the `Makefile`s when the user runs `./configure` before building.

GNU Autoconf (`autoconf`)
: GNU Autoconf will build the `configure` script using the configurations we have defined (hand-written) in `configure.ac`.

GNU Autoconf Archive
: These are a large collection of tests that can be called to run at `./configure` time. See the explanation under GNU Portability Library above for instructions on obtaining it and keeping it up to date.

GNU Libtool (`libtool`)
: GNU Libtool is in charge of building all the libraries in Gnuastro. The libraries contain functions that are used by more than one program and are installed for use in other programs. They are thus put in a separate directory (`lib/`).

GNU help2man (`help2man`)
: GNU help2man is used to convert the output of the `--help` option (Section 4.7.2 [`--help`], page 75) to the traditional Man page (Section 4.7.3 [Man pages], page 76).

LaTeX and some TeX packages
: Some of the figures in this book are built by LaTeX (using the PGF/TikZ package). The LaTeX source for those figures is version controlled for easy maintenance not the actual figures. So the `./boostrap` script will run LaTeX to build the figures. The best way to install LaTeX and all the necessary packages is

---

[5] If you are not a developer in Gnulib or Autoconf archives, `DEVDIR` can be a directory that you don't backup. In this way the large number of files in these projects won't slow down your backup process or take bandwidth (if you backup to a remote server).

[6] If you already have the Autoconf archives in a separate directory, or can't clone it in the same directory as Gnulib, or you have it with another directory name (not `autoconf-archive/`), you can follow this short step. Set `AUTOCONFARCHIVES` to your desired address. Then define a symbolic link in `DEVDIR` with the following command so Gnuastro's bootstrap script can find it:
`$ ln -s $AUTOCONFARCHIVES $DEVDIR/autoconf-archive`.

[7] If your internet connection is active, but Git complains about the network, it might be due to your network setup not recognizing the git protocol. In that case use the following URL for the HTTP protocol instead (for Autoconf archives, replace the name): `http://git.sv.gnu.org/r/gnulib.git`

through TEX live (`https://www.tug.org/texlive/`) which is a package manager for TEX related tools that is independent of any operating system. It is thus preferred to the TEX Live versions distributed by your operating system.

To install TEX Live, go to the webpage and download the appropriate installer by following the "download" link. Note that by default the full package repository will be downloaded and installed (around 4 Giga Bytes) which can take *very* long to download and to update later. However, most packages are not needed by everyone, it is easier, faster and better to install only the "Basic scheme" (consisting of only the most basic TEX and LATEX packages, which is less than 200 Mega bytes)[8].

After the installation be sure to set the environment variables as suggested in the end of the outputs. Any time you confront (need) a package you don't have, simply install it with a command like below (similar to how you install software from your operating system's package manager)[9]. To install all the necessary TEX packages for a successful Gnuastro bootstrap, run this command:

```
$ su
# tlmgr install epsf jknapltx caption biblatex biber iftex \
                 etoolbox logreq xstring xkeyval pgf ms     \
                 xcolor pgfplots times rsfs pstools epspdf
```

ImageMagick (`imagemagick`)

ImageMagick is a wonderful and robust program for image manipulation on the command-line. `bootsrap` uses it to convert the book images into the formats necessary for the various book formats.

3.2 Downloading the source

Gnuastro's source code can be downloaded in two ways. As a tarball, ready to be configured and installed on your system (as described in Section 1.1 [Quick start], page 1), see Section 3.2.1 [Release tarball], page 31. If you want official releases of stable versions this is the best, easiest and most common option. Alternatively, you can clone the version controlled history of Gnuastro, run one extra bootstrapping step and then follow the same steps as the tarball. This will give you access to all the most recent work that will be included in the next release along with the full project history. The process is thoroughly introduced in Section 3.2.2 [Version controlled source], page 32.

3.2.1 Release tarball

A release tarball (commonly compressed) is the most common way of obtaining free and open source software. A tarball is a snapshot of one particular moment in the Gnuastro development history along with all the necessary files to configure, build, and install Gnuastro easily (see Section 1.1 [Quick start], page 1). It is very straightforward and needs the least set of dependencies (see Section 3.1.1 [Mandatory dependencies], page 26). Gnuastro has tarballs for official stable releases and pre-releases for testing. See Section 1.5 [Version

[8] You can also download the DVD iso file at a later time to keep as a backup for when you don't have internet connection if you need a package.

[9] After running TEX, or LATEX, you might get a warning complaining about a `missingfile`. Run '`tlmgr info missingfile`' to see the package(s) containing that file which you can install.

numbering], page 5, for more on the two types of releases and the formats of the version numbers. The URLs for each type of release are given below.

Official stable releases (`http://ftp.gnu.org/gnu/gnuastro`):
: This URL hosts the official stable releases of Gnuastro. Always use the most recent version (see Section 1.5 [Version numbering], page 5). By clicking on the "Last modified" title of the second column, the files will be sorted by their date which you can also use to find the latest version. It is recommended to use a mirror to download these tarballs, please visit `http://ftpmirror.gnu.org/gnuastro/` and see below.

Pre-release tar-balls (`http://alpha.gnu.org/gnu/gnuastro`):
: This URL contains unofficial pre-release versions of Gnuastro. The pre-release versions of Gnuastro here are for enthusiasts to try out before an official release. If there are problems, or bugs then the testers will inform the developers to fix before the next official release. See Section 1.5 [Version numbering], page 5, to understand how the version numbers here are formatted. If you want to remain even more up-to-date with the developing activities, please clone the version controlled source as described in Section 3.2.2 [Version controlled source], page 32.

Gnuastro's official tarball is released with two formats: Gzip (with suffix `.tar.gz`) and Lzip (with suffix `.tar.lz`). The former is a very well-known and widely used compression program created by GNU and available in most systems. The latter provides a better compression ratio and more robust archival capacity. For example Gnuastro 0.2's tarball was 2.8MB and 4.2MB with Lzip and Gzip respectively. From the Lzip webpage (`http://www.nongnu.org/lzip/lzip.html`): "Lzip is a lossless data compressor with a user interface similar to the one of gzip or bzip2. Lzip can compress about as fast as gzip (`lzip -0`), or compress most files more than bzip2 (`lzip -9`). Decompression speed is intermediate between gzip and bzip2. Lzip is better than gzip and bzip2 from a data recovery perspective." However, Lzip is currently not too common, so it might not be pre-installed in your operating system. Installing it from your operating system's package manager or from source is very easy.

The GNU FTP server is mirrored (has backups) in various locations on the globe (`http://www.gnu.org/order/ftp.html`). You can use the closest mirror to your location for a more faster download. Note that only some mirrors keep track of the pre-release (alpha) tarballs. Also note that if you want to download immediately after and announcement (see Section 1.9 [Announcements], page 11), the mirrors might need some time to synchronize with the main GNU FTP server.

3.2.2 Version controlled source

The publicly distributed Gnuastro tar-ball (for example `gnuastro-X.X.tar.gz`) does not contain the revision history, it is only a snapshot of the source code at one significant instant of Gnuastro's history (specified by the version number, see Section 1.5 [Version numbering], page 5), ready to be configured and built. To be able to develop successfully, the revision history of the code can be very useful to track when something was added or changed, also some updates that are not yet officially released might be in it.

We use Git for the version control of Gnuastro. For those who are not familiar with it, we recommend the Pro Git[10] book. The whole book is publicly available for online reading and downloading and does a wonderful job at explaining the concepts and best practices.

Let's assume you want to keep Gnuastro in the `TOPGNUASTRO` directory (can be any directory, change the value below). The full version controlled history of Gnuastro can be cloned in `TOPGNUASTRO/gnuastro` by running the following commands[11]:

```
$ TOPGNUASTRO=/home/yourname/Research/projects/
$ cd $TOPGNUASTRO
$ git clone git://git.sv.gnu.org/gnuastro.git
```

The `$TOPGNUASTRO/gnuastro` directory will contain hand-written (version controlled) source code for Gnuastro's programs, libraries, this book and the tests. All are divided into sub-directories with standard and very descriptive names. The version controlled files in the top cloned directory are either mainly in capital letters (for example `THANKS` and `README`) or mainly written in small-caps (for example `configure.ac` and `Makefile.am`). The former are non-programming, standard writing for human readers containing high-level information about the whole package. The latter are instructions to customize the GNU build system for Gnuastro.

The cloned Gnuastro source cannot immediately be configured, compiled, or installed since it only contains hand-written files, not automatically generated or imported files which do all the hard work of the build process. See Section 3.2.2.1 [Bootstrapping], page 33, for the process of generating and importing those files (it is very easy!). Once you have bootstrapped Gnuastro, you can run the standard procedures (in Section 1.1 [Quick start], page 1). Very soon after you have cloned it, Gnuastro's main `master` branch will be updated on the main repository (since the developers are actively working on Gnuastro), for the best practices in keeping your local history in sync with the main repository see Section 3.2.2.2 [Synchronizing], page 35.

3.2.2.1 Bootstrapping

The version controlled source code lacks the source files that we have not written or are automatically built. These automatically generated files are included in the distributed tar ball for each distribution (for example `gnuastro-X.X.tar.gz`, see Section 1.5 [Version numbering], page 5) and make it easy to immediately configure, build, and install Gnuastro. However from the perspective of version control, they are just bloatware and sources of confusion (since they are not changed by Gnuastro developers).

The process of automatically building and importing necessary files into the cloned directory is known as *bootstrapping*. All the instructions for an automatic bootstrapping are available in `bootstrap` and configured using `bootstrap.conf`. `bootstrap` is the only file not written by Gnuastro developers but is under version control to enable simple bootstrapping immediately after cloning. It is maintained by the GNU Portability Library (Gnulib) and this file is an identical copy, so do not make any changes in this file since it will be replaced when Gnulib releases an update. Make all your changes in `bootstrap.conf`.

---

[10] `https://progit.org/`

[11] If your internet connection is active, but Git complains about the network, it might be due to your network setup not recognizing the Git protocol. In that case use the following URL which uses the HTTP protocol instead: `http://git.sv.gnu.org/r/gnuastro.git`

The bootstrapping process has its own separate set of dependencies, the full list is given in Section 3.1.3 [Bootstrapping dependencies], page 29. They are generally very low-level and used by a very large set of commonly used programs, so they are probably already installed on your system. The simplest way to bootstrap Gnuastro is to simply run the bootstrap script within your cloned Gnuastro directory as shown below. However, please read the next paragraph before doing so (see Section 3.2.2 [Version controlled source], page 32, for `TOPGNUASTRO`).

```
$ cd TOPGNUASTRO/gnuastro
$ ./bootstrap                      # Requires internet connection
```

Without any options, `bootstrap` will clone Gnulib within your cloned Gnuastro directory (`TOPGNUASTRO/gnuastro/gnulib`) and download the necessary Autoconf archives macros. So if you run bootstrap like this, you will need an internet connection every time you decide to bootstrap. Also, Gnulib is a large package and cloning it can be slow. It will also keep the full Gnulib repository within your Gnuastro repository, so if another one of your projects also needs Gnulib, and you insist on running bootstrap like this, you will have two copies. In case you regularly backup your important files, Gnulib will also slow down the backup process. Therefore while the simple invocation above can be used with no problem, it is not recommended. To do better, see the next paragraph.

The recommended way to get these two packages is thoroughly discussed in Section 3.1.3 [Bootstrapping dependencies], page 29, (in short: clone them in the separate `DEVDIR/` directory). The following commands will take you into the cloned Gnuastro directory and run the `bootstrap` script, while telling it to copy some files (instead of making symbolic links, with the `--copy` option, this is not mandatory[12]) and where to look for Gnulib (with the `--gnulib-srcdir` option).

```
$ cd $TOPGNUASTRO/gnuastro
$ ./bootstrap --copy --gnulib-srcdir=$DEVDIR/gnulib
```

Since Gnulib and Autoconf archives are now available in your local directories, you don't need an internet connection every time you decide to remove all untracked files and redo the bootstrap (see box below). You can also use the same command on any other project that uses Gnulib. All the necessary GNU C library functions, Autoconf macros and Automake inputs are now available along with the book figures. The standard GNU build system (Section 1.1 [Quick start], page 1) will do the rest of the job.

Undoing the bootstrap: During the development, it might happen that you want to remove all the automatically generated and imported files. In other words, you might want to reverse the bootstrap process. Fortunately Git has a good program for this job: `git clean`. Run the following command and every file that is not version controlled will be removed.

```
git clean -fxd
```

It is best to commit any recent change before running this command. You might have created new files since the last commit and if they haven't been committed, they will all be gone forever (using `rm`). To get a list of the non-version controlled files instead of deleting them, add the `n` option to `git clean`, so it becomes `-fxdn`.

---

[12] The `--copy` option is recommended because some backup systems might do strange things with symbolic links.

Besides the `bootstrap` and `bootstrap.conf`, the `bootstrapped/` directory and `README-hacking` file are also related to the bootstrapping process. The former hosts all the imported (bootstrapped) directories. Thus, in the version controlled source, it only contains a `REAME` file, but in the distributed tar-ball it also contains sub-directories filled with all bootstrapped files. `README-hacking` contains a summary of the bootstrapping process discussed in this section. It is a necessary reference when you haven't built this book yet. It is thus not distributed in the Gnuastro tarball.

3.2.2.2 Synchronizing

The bootstrapping script (see Section 3.2.2.1 [Bootstrapping], page 33) is not regularly needed: you mainly need it after you have cloned Gnuastro (once) and whenever you want to re-import the files from Gnulib, or Autoconf archives[13] (not too common). However, Gnuastro developers are constantly working on Gnuastro and are pushing their changes to the official repository. Therefore, your local Gnuastro clone will soon be out-dated. Gnuastro has two mailing lists dedicated to its developing activities (see Section 11.10 [Developing mailing lists], page 335). Subscribing to them can help you decide when to synchronize with the official repository.

To pull all the most recent work in Gnuastro, run the following command from the top Gnuastro directory:

```
$ git pull && autoconf -f
```

GNU Autoconf is part of the GNU build system and will update the `./configure` script based on the hand-written configurations (in `configure.ac`, which is version controlled in Gnuastro). The pulled changes might contain changes in the build system configurations. However, The most important reason for running this command is to generate a version number for your Gnuastro snapshot. This generated version number will include the commit information if you are building Gnuastro from any point in Gnuastro's history (see Section 1.5 [Version numbering], page 5). Since the version number is included in nearly all outputs of the programs, this can help you later exactly reproduce an old result by checking out the exact point in Gnuastro's history that produced those results. Therefore, be sure to run '`autoconf -f`' after every synchronization. You can also run them separately:

```
$ git pull
$ autoconf -f
```

If you would like to see what has changed since you last synchronized your local clone, you can take the following steps instead of the simple command above (don't type anything after `#`):

```
$ git checkout master              # Confirm if you are on master.
$ git fetch origin                 # Fetch all new commits from server.
$ git log master..origin/master    # See all the new commit messages.
$ git merge origin/master          # Update your master branch.
$ autoconf -f                      # Update ./configure.
```

By default `git log` prints the most recent commit first, add the `--reverse` option to see the changes chronologically. To see exactly what has been changed in the source code along with the commit message, add a `-p` option to the `git log`.

---

[13] `https://savannah.gnu.org/task/index.php?13993` is defined for you to check if significant (for Gnuastro) updates are made in these repositories, since the last time you pulled from them.

If you intend make changes in the code, have a look at Chapter 11 [Developing], page 320, to get started easily. Be sure to commit your changes in a separate branch (keep your `master` branch to follow the official repository) and re-run `autoconf -f` after the commit. If you intend to send your changes to us (see Section 11.11 [Contributing to Gnuastro], page 335) for the benefit of the whole community. If you send your work to us, you can safely use your commit since it will be ultimately recorded in Gnuastro's official history. If not, please upload your separate branch to a public hosting service (for example GitLab, see Section 11.11.4 [Forking tutorial], page 339) and link to it in your report, or run `make distcheck` and upload the output `gnuastro-X.X.X.XXXX.tar.gz` to a publicly accessible webpage so your results can be considered scientific (reproducible).

3.3 Build and install

This section is basically a longer explanation to the sequence of commands given in Section 1.1 [Quick start], page 1. If you didn't have any problems during the Section 1.1 [Quick start], page 1, steps, you want to have all the programs of Gnuastro installed in your system, you don't want to change the executable names during or after installation, you have root access to install the programs in the default system wide directory, the Letter paper size of the print book is fine for you or as a summary you don't feel like going into the details when everything is working, you can safely skip this section.

If you have any of the above problems or you want to understand the details for a better control over your build and install, read along. The dependencies which you will need prior to configuring, building and installing Gnuastro are explained in Section 3.1 [Dependencies], page 25. The first three steps in Section 1.1 [Quick start], page 1, need no extra explanation, so we will skip them and start with an explanation of Gnuastro specific configuration options and a discussion on the installation directory in Section 3.3.1 [Configuring], page 36, followed by some smaller subsections: Section 3.3.2 [Tests], page 45, Section 3.3.3 [A4 print book], page 45, and Section 3.3.4 [Known issues], page 46, which explains the solutions to known problems you might encounter in the installation steps and ways you can solve them.

3.3.1 Configuring

The `$ ./configure` step is the most important step in the build and install process. All the required packages, libraries, headers and environment variables are checked in this step. The behaviors of make and make install can also be set through command line options to this command.

The configure script accepts various arguments and options which enable the final user to highly customize whatever she is building. The options to configure are generally very similar to normal program options explained in Section 4.1.1 [Arguments and options], page 48. Similar to all GNU programs, you can get a full list of the options along with a short explanation by running

```
$ ./configure --help
```

A complete explanation is also included in the `INSTALL` file. Note that this file was written by the authors of GNU Autoconf (which builds the `configure` script), therefore it is common for all programs which use the `$ ./configure` script for building and installing, not just Gnuastro. Here we only discuss cases where you don't have super-user access to the system

and if you want to change the executable names. But before that, a review of the options to configure that are particular to Gnuastro are discussed.

3.3.1.1 Gnuastro configure options

Most of the options to configure (which are to do with building) are similar for every program which uses this script. Here the options that are particular to Gnuastro are discussed. The next topics explain the usage of other configure options which can be applied to any program using the GNU build system (through the configure script).

`--enable-progname`
: Only build and install `progname` along with any other program that is enabled in this fashion. `progname` is the name of the executable without the `ast`, for example `crop` for Crop (with the executable name of `astcrop`). If this option is called for any of the programs in Gnuastro, any program which is not explicitly enabled will not be built or installed.

`--disable-progname`
`--enable-progname=no`
: Do not build or install the program named `progname`. This is very similar to the `--enable-progname`, but will build and install all the other programs except this one.

`--enable-bin-op-uint8`
`--enable-bin-op-int8`
`--enable-bin-op-uint16`
`--enable-bin-op-int16`
`--enable-bin-op-uint32`
`--enable-bin-op-int32`
`--enable-bin-op-uint64`
`--enable-bin-op-int64`
`--enable-bin-op-float32`
`--enable-bin-op-float64`
: Enable the binary data-structure operators to work natively on the respective type of data (`u` stands for unsigned types, see Section 4.4 [Numeric data types], page 64). Some are compiled by default, to disable them (or disable any other type), either run `enable-bin-op-TYPE=no`, or run `--disable-bin-op-TYPE`. The final list of enabled/disabled types can be inspected in the outputs of `./configure` (close to the end).

  Binary operators, for example `+` or `>` (greater than), are some of the most common operators to the Section 6.2 [Arithmetic], page 108, program or the `data_arithmetic` function in Section 10.3 [Gnuastro library], page 233. To operate most efficiently (as fast as possible without using extra memory or CPU resources), it is best to rely on the native types of the input data. For example, if you want to add an integer array with a floating point array, using the native types, means relying the system's internal type conversion for each array element, see Section 6.2.3 [Invoking Arithmetic], page 114. If we don't use the native conversion, then the integer array has to be converted to the same type as the floating point array to do the conversion. This will consume memory

(to copy the integer array into a new float array) and CPU (integer types need much less processing) resources and ultimately slow down the running.

There are many binary operators and in order to have them operate natively on of each of the above types, the compiler has to prepare for all the different combinations of these types. This can greatly slow down the compilation[14] (when you run `make`). For example, with only one type, `make` will finish in less than a minute, but if you enable all types, it can take roughly half an hour. However, the profits of this one-time investment at compilation time will be directly felt (more significantly on large images/datasets) each time you run Gnuastro programs or libraries, because no internal type conversion will be necessary.

If build time is important for you (mainly developers), disabling shared libraries and optimizations (as in Section 11.6 [Building and debugging], page 331) is the first step to take. If you commonly work with very specific data-types, you can enable them (and disable the default types that you don't need) with these configuration options. Since the outputs of comparison operators are `unsigned char` (or `uint8_t`) type and most astronomical datasets are in single precision (32-bit) floating point (`float`), the recommended minimum enabled types are `uint8` and `float32`.

GNU/Linux distribution package managers who compile once, for a large audience of users who just download the compiled programs and executables, are recommended to enable all types to help their users.

`--enable-bin-op-alltypes`
: Enable native binary arithmetic operation on all types, see the description above for the various types for a full discussion. As discussed there, enabling all types can greatly speed up arithmetic operations on any arbitrary dataset, but will also slow down the building time of Gnuastro. Recall that in practice this only affects the Section 6.2 [Arithmetic], page 108, program and the `gal_arithmetic` library binary operators, nothing else. This option is strongly recommended when you are building Gnuastro to be included in a package manager of a GNU/Linux distribution (or other operating system).

`--enable-gnulibcheck`
: Enable checks on the GNU Portability Library (Gnulib). Gnulib is used by Gnuastro to enable users of non-GNU based operating systems (that don't use GNU C library or glibc) to compile and use the advanced features that this library provides. We make extensive use of such functions. If you give this option to `$ ./configure`, when you run `$ make check`, first the functions in Gnulib will be tested, then the Gnuastro executables. If your operating system does not support glibc or has an older version of it and you have problems in the build process (`$ make`), you can give this flag to configure to see if the problem is caused by Gnulib not supporting your operating system or Gnuastro, see Section 3.3.4 [Known issues], page 46.

[14] It can also greatly increase the file size of the library, from a few hundred kilobytes to a few megabytes.

`--disable-guide-message`
`--enable-guide-message=no`

Do not print a guiding message during the GNU Build process of Section 1.1 [Quick start], page 1. By default, after each step, a message is printed guiding the user what the next command should be. Therefore, after `./configure`, it will suggest running `make`. After `make`, it will suggest running `make check` and so on. If Gnuastro is configured with this option, for example

```
$ ./configure --disable-guide-message
```

Then these messages will not be printed after any step (like most programs). For people who are not yet fully accustomed to this build system, these guidelines can be very useful and encouraging. However, if you find those messages annoying, use this option.

Note: If some programs are enabled and some are disabled, it is equivalent to simply enabling those that were enabled. Listing the disabled programs is redundant.

Note that the tests of some programs might require other programs to have been installed and tested. For example MakeProfiles is the first program to be tested when you run `$ make check`, it provides the inputs to all the other tests. So if you don't install MakeProfiles, then the tests for all the other programs will be skipped or fail. To avoid this, in one run, you can install all the packages and run the tests but not install. If everything is working correctly, you can run configure again with only the packages you want but not run the tests and directly install after building.

3.3.1.2 Installation directory

One of the most commonly used options to configure is the directory that will host all the files which require installing (for example the actual executable files for the program and/or the documentation and configuration files). This is done through the `--prefix` option. To demonstrate its applicability, let's assume you don't have administrator or root access but you want to be able to access the installed files from anywhere while you are logged in (one of the most common uses of `--prefix`).

The most basic way to run an executable is to explicitly type the full file name (including all the directory information) and run it. One example is running the configuration script with the `$ ./configure` command (see Section 1.1 [Quick start], page 1). By giving a specific directory (the current directory or `./`), we are explicitly telling the shell to look in the current directory for an executable file named '`configure`'. Directly specifying the directory is thus useful for simple shell scripts in the current (or nearby) directories. However, when the program (an executable) is to be used a lot, specifying all those directories will become a significant burden. For example, the `ls` executable lists the contents in a given directory and it is (usually) installed in the `/usr/bin/` directory by the operating system maintainers. So each time you you want to use it you would have to run the following command (which is very inconvenient, both in typing and in remembering the various directories).

```
$ /usr/bin/ls
```

To address this problem, we have the `PATH` environment variable. To understand it better, we will start with a short introduction to the shell variables. Shell variable values

are basically treated as strings of characters. You can define a variable and a value for it by running

```
$ myvariable=a_test_value
```

You can see the value it represents by running

```
$ echo $myvariable
```

If a variable has no value, the last command will only print an empty line. A variable defined like this will be known as long as this shell or terminal is running. Other terminals will have no idea it existed. The main advantage of shell variables is that if they are exported[15], subsequent programs that are run within that shell can access their value. So by giving them different values, you can change the "environment" of a program which uses their values. The shell variables which are accessed by programs are therefore known as "environment variables"[16]. You can see the full list of the environment variables that your shell currently recognizes by running:

```
$ printenv
```

`HOME` is one commonly used environment variable, it is any user's (the one that is logged in) top directory. It is used so often that the shell has a special expansion (alternative) for it: '`~`'. Whenever you see file names starting with the tilde sign, it actually represents the value to the `HOME` environment variable, so `~/doc` is the same as `$HOME/doc`.

Another one of the most commonly used environment variables is `PATH`, it is a list of directories to search for executable names. Its value is a list of directories (separated by a colon, or '`:`'). When the address of the executable is not explicitly given (like `./configure` above), the system will look for the executable in the directories specified by `PATH`. If you have a computer nearby, try running the following command to see which directories your system will look into when it is searching for executable (binary) files, one example is printed here (notice how `/usr/bin`, in the `ls` example above, is one of the directories in `PATH`):

```
$ echo $PATH
/usr/local/sbin:/usr/local/bin:/usr/bin
```

By default `PATH` usually contains system-wide directories, which are readable (but not writable) by all users, like the above example. Therefore if you don't have root (or administrator) access, you need to add another directory to `PATH` which you actually have write access to. The standard directory where you can keep installed files (not just executables) for your own user is the `~/.local/` directory. The names of hidden files start with a '`.`' (dot), so it will not show up in your common command-line listings, or on the graphical user interface. You can use any other directory, but this is the most recognized.

The top installation directory will be used to keep all the package's components: programs (executables), libraries, include (header) files, shared data (like manuals), or configuration files (see Section 10.1 [Review of library fundamentals], page 223, for a thorough introduction to headers and linking). So it commonly has some of the following sub-directories for each class of components respectively: `bin/`, `lib/`, `include/` `man/`, `share/`, `etc/`. Since the `PATH` variable is only used for executables, you can add the `~/.local/bin` directory to `PATH` with the following command. As defined below, first the existing value of `PATH` is

---

[15] By running `$ export myvariable=a_test_value` instead of the simpler case in the text

[16] You can use shell variables for other actions too, for example to temporarily keep some names or run loops on some files.

used, then your given directory is added to its end and the combined value is put back in PATH (run '$ echo $PATH' afterwards to check if it was added).

```
$ PATH=$PATH:~/.local/bin
```

Any executable that you installed in this directory will now be usable without having to remember and type its full address. However, as soon as you leave your current terminal session, this modified PATH variable will be forgotten. Adding your specified directory to the PATH environment variable each time you start a terminal is also very inconvenient and prone to errors. So there are standard 'startup files' defined by your shell. There is a special startup file for every significant starting step:

/etc/profile and everything in /etc/profile.d/
: These startup scripts are called when your whole system starts (for example after you turn on your computer). Therefore you need administrator or root privileges to access or modify them.

~/.bash_profile
: If you are using (GNU) Bash as your shell, the commands in this file are run once every time you log in to your account.

~/.bashrc
: If you are using (GNU) Bash as your shell, the commands here will be run each time you start a terminal (for example, when you open your terminal emulator in the graphic user interface).

For security reasons, it is highly recommended to directly type in your HOME directory value by hand in startup files instead of using variables. So in the following, let's assume your user name is 'name'. To add ~/.local/bin to your PATH automatically on any startup file, you have to "export" the new value of PATH in the startup file that is most relevant to you with the line 'export PATH=$PATH:/home/name/.local/bin'. You can either do it manually using a text editor, or by running the following command which will add this line as the last line of the file. Let's assume you want to add it to ~/.bashrc (afterwards, open your ~/.bashrc with a text editor and check it out, to see the result for your self):

```
$ echo 'export PATH=$PATH:/home/name/.local/bin' >> ~/.bashrc
```

Now that you know your system will look into ~/.local/bin for executables, you can tell Gnuastro's configure script to install everything in the top ~/.local directory using the --prefix option. The configure script will then put the executables in ~/.local/bin, the compiled library files in ~/.local/lib, the library header files in ~/.local/include and so on (see Section 10.1 [Review of library fundamentals], page 223, for a review of the compiled library files and the headers). When you subsequently run $ make install, all the install-able files will be put in their respective directory under ~/.local/ (as discussed above). Note that tilde ('~') expansion will not happen if you put a '=' between --prefix and ~/.local[17], so we have avoided the = character here which is optional in GNU-style options, see Section 4.1.1.2 [Options], page 50.

```
$ ./configure --prefix ~/.local
```

You can install everything (including libraries like GSL, CFITSIO, or WCSLIB) locally by configuring them as above. However, recall that PATH is only for executable files, not

[17] If you insist on using '=', you can use --prefix=$HOME/.local.

libraries. Therefore, when building programs or libraries[18] that depend on libraries you installed like this, you have to guide your compiler to the necessary directories. To do that, you have to define the `LDFLAGS` and `CPPFLAGS` environment variables respectively. This can be done while calling `./configure` as shown below:

```
$ ./configure LDFLAGS=-L/home/name/.local/lib            \
              CPPFLAGS=-I/home/name/.local/include       \
              --prefix ~/.local
```

If you do this for all your libraries, it can be annoying to do this when configuring every software that depends on them. Hence, you can define these two variables in the most relevant startup file (discussed above). The convention on using these variables doesn't include a colon to separate values (as `PATH`-like variables do), they use white space characters and each value is prefixed with a compiler option[19] (note the `-L` and `-I` above, see Section 4.1.1.2 [Options], page 50, Section 10.1.1 [Headers], page 224, for `-I`, and Section 10.1.2 [Linking], page 227, for `-L`). Therefore we have to keep the value in double quotation signs to keep the white space characters. Run the following two commands to do that if you want them in your `~/.bashrc`.

```
echo 'export LDFLAGS="$LDFLAGS -L/home/name/.local/lib"' >> ~/.bashrc
echo 'export CPPFLAGS="$CPPFLAGS -I/home/name/.local/include"' \
     >> ~/.bashrc
```

Dynamic libraries are linked to the executable every time you run a program that depends on them (see Section 10.1.2 [Linking], page 227, to fully understand this important concept). Hence dynamic libraries also require a special path variable called `LD_LIBRARY_PATH`. To use programs that depend on these libraries, you will need to have `~/.local/lib` added to your `LD_LIBRARY_PATH` environment variable in the relevant start-up file (similar to `PATH`) as shown below (the \ is only to fit this command within the page width, you can ignore it when typing on the terminal):

```
echo 'export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/home/name/.local/lib' \
     >> ~/.bashrc
```

If you also want to access the Info (see Section 4.7.4 [Info], page 76) and man pages (see Section 4.7.3 [Man pages], page 76) documentations add `~/.local/share/info` and `~/.local/share/man` to your `INFOPATH`[20] and `MANPATH` environment variables respectively.

A final note is that order matters in the directories that are searched for all the variables discussed above. In the examples above, the new directory was added after the system specified directories. So if the program, library or manuals are found in the system wide directories, the user directory is no longer searched. If you want to search your local

[18] For example WCSLIB which needs CFITSIO, or Gnuastro which needs both.

[19] These variables are ultimately used as options while building the programs, so every value has be an option name followed be a value as discussed in Section 4.1.1.2 [Options], page 50.

[20] Info has the following convention: "If the value of `INFOPATH` ends with a colon [or it isn't defined] ..., the initial list of directories is constructed by appending the build-time default to the value of `INFOPATH`." So when installing in a non-standard directory and if `INFOPATH` was not initially defined, add a colon to the end of `INFOPATH` as shown below, otherwise Info will not be able to find system-wide installed documentation:

```
echo 'export INFOPATH=$INFOPATH:/home/name/.local/share/info:' >> ~/.bashrc
```

Note that this is only an internal convention of Info, do not use it for other `*PATH` variables.

installation first, put the new directory before the already existing list, like the example below.

```
export LD_LIBRARY_PATH=/home/name/.local/lib:$LD_LIBRARY_PATH
```

This is good when a library, for example CFITSIO, is already present on the system, but the system-wide install wasn't configured with the correct configuration flags (see Section 3.1.1.2 [CFITSIO], page 26), or you want to use a newer version and you don't have administrator or root access to update it. If you update `LD_LIBRARY_PATH` with the `~/.local/lib` installation first, the linker will first find the CFITSIO you installed for yourself and link with it. It will never reach the system-wide installation. However there are important security problems with this order: all important system-wide programs and libraries can be replaced by non-secure versions with the same file names and put in the customized directory (`./.local/lib` in this example). So if you choose to search in your customized directory first, be sure to keep it clean from executables or libraries with the same names as important system programs or libraries.

3.3.1.3 Executable names

At first sight, the names of the executables for each program might seem to be uncommonly long, for example `astnoisechisel` or `astcrop`. We could have chosen terse (and cryptic) names like most programs do. We chose this complete naming convention (something like the commands in TeX) so you don't have to spend too much time remembering what the name of a specific program was. Such complete names also enable you to easily search for the programs.

To facilitate typing the names in, we suggest using the shell auto-complete. With this facility you can find the executable you want very easily. It is very similar to file name completion in the shell. For example, simply by typing the letters below (where `[TAB]` stands for the Tab key on your keyboard)

```
$ ast[TAB][TAB]
```

you will get the list of all the available executables that start with `ast` in your `PATH` environment variable directories. So, all the Gnuastro executables installed on your system will be listed. Typing the next letter for the specific program you want along with a Tab, will limit this list until you get to your desired program.

In case all of this does not convince you and you still want to type short names, some suggestions are given below. You should have in mind though, that if you are writing a shell script that you might want to pass on to others, it is best to use the standard name because other users might not have adopted the same customizations. The long names also serve as a form of documentation in such scripts. A similar reasoning can be given for option names in scripts: it is good practice to always use the long formats of the options in shell scripts, see Section 4.1.1.2 [Options], page 50.

The simplest solution is making a symbolic link to the actual executable. For example let's assume you want to type `ic` to run Crop instead of `astcrop`. Assuming you installed Gnuastro executables in `/usr/local/bin` (default) you can do this simply by running the following command as root:

```
# ln -s /usr/local/bin/astcrop /usr/local/bin/ic
```

In case you update Gnuastro and a new version of Crop is installed, the default executable name is the same, so your custom symbolic link still works.

The installed executable names can also be set using options to `$ ./configure`, see Section 3.3.1 [Configuring], page 36. GNU Autoconf (which configures Gnuastro for your particular system), allows the builder to change the name of programs with the three options `--program-prefix`, `--program-suffix` and `--program-transform-name`. The first two are for adding a fixed prefix or suffix to all the programs that will be installed. This will actually make all the names longer! You can use it to add versions of program names to the programs in order to simultaneously have two executable versions of a program.

The third configure option allows you to set the executable name at install time using the SED program. SED is a very useful 'stream editor'. There are various resources on the internet to use it effectively. However, we should caution that using configure options will change the actual executable name of the installed program and on every re-install (an update for example), you have to also add this option to keep the old executable name updated. Also note that the documentation or configuration files do not change from their standard names either.

For example, let's assume that typing `ast` on every invocation of every program is really annoying you! You can remove this prefix from all the executables at configure time by adding this option:

```
$ ./configure --program-transform-name='s/ast/ /'
```

3.3.1.4 Configure and build in RAM

The configure and build process involves the creation, reading, and modification of a large number of files (input/output, or I/O). Therefore file I/O issues can directly affect the work of developers who need to configure and build Gnuastro numerous times. Some of these issues are listed below:

- I/O will cause wear and tear on both the HDDs (mechanical failures) and SSDs (decreasing the lifetime).
- Having the built files mixed with the source files can greatly affect backing up (synchronization) of source files (since it involves the management of a large number of small files that are regularly changed. Backup software can of course be configured to ignore the built files and directories. However, since the built files are mixed with the source files and can have a large variety, this will require a high level of customization.

One solution to address both these problems is to use the tmpfs file system (`https://en.wikipedia.org/wiki/Tmpfs`). Any file in tmpfs is actually stored in the RAM (and possibly SAWP), not on HDDs or SSDs. The RAM is built for extensive and fast I/O. Therefore the large number of file I/Os associated with configuring and building will not harm the HDDs or SSDs. Due to the volatile nature of RAM, files in the tmpfs file-system will be permanently lost after a power-off. Since all configured and built files are derivative files (not files that have been directly written by hand) there is no problem in this and this feature can be considered as an automatic cleanup.

The modern GNU C library (and thus the Linux kernel) define the `/dev/shm` directory for this purpose (POSIX shared memory). So using GNU Build System's ability to build in a separate directory (not necessarily in the source directory), we can configure and build the programs in `/dev/shm` to benefit from the RAM. To simplify the process, Gnuastro comes with a `tmpfs-config-make` script. This script will create a directory in the shared memory, and put a symbolic link to it (called `build`) in the top source directory (the backup/sync

software therefore only needs to ignore this single link/file). The script will then internally change to that directory and configure and build (`make -kjN`, where `N` is the number of threads for a parallel build) Gnuastro in there. To benefit from this script, simply run the following command instead of `./configure` and `make`:

```
$ ./tmpfs-config-make
```

After this script is finished, you can '`cd build`' and run other Make commands (for example, '`make check`', '`make install`', or '`make pdf`') from there. In Emacs, the command to be run with the `M-x compile` command (by default: `make -k`) can be changed to '`cd build; make -kjN`', or '`make -C build -kjN`' (`N` is the number of threads; an integer ≥ 1). For subsequent builds (during development) the `M-x recompile` command will also do all the building in the RAM while you modify the clean, and backed-up source files and make minimal/efficient use of your non-volatile HDD or SSD.

This script can be used in any software which is configured and built using the GNU Build System. Just copy it in the top source directory of that software and run it from there. The default number of threads and location of the shared memory (`/dev/shm`) are currently hard-coded within the script. If you need to change them, please open the script with a text editor and set their values manually.

3.3.2 Tests

After successfully building (compiling) the programs with the `$ make` command you can check the installation before installing. To run the tests, run

```
$ make check
```

For every program some tests are designed to check some possible operations. Running the command above will run those tests and give you a final report. If everything is ok and you have built all the programs, all the tests should pass. In case any of the tests fail, please have a look at Section 3.3.4 [Known issues], page 46, and if that still doesn't fix your problem, look that the `./tests/test-suite.log` file to see if the source of the error is something particular to your system or more general. If you feel it is general, please contact us because it might be a bug. Note that the tests of some programs depend on the outputs of other program's tests, so if you have not installed them they might be skipped or fail. Prior to releasing every distribution all these tests are checked. If you have a reasonably modern terminal, the outputs of the successful tests will be colored green and the failed ones will be colored red.

These scripts can also act as a good set of examples for you to see how the programs are run. All the tests are in the `tests/` directory. The tests for each program are shell scripts (ending with `.sh`) in a sub-directory of this directory with the same name as the program. See Section 11.7 [Test scripts], page 332, for more detailed information about these scripts in case you want to inspect them.

3.3.3 A4 print book

The default print version of this book is provided in the letter paper size. If you would like to have the print version of this book on paper and you are living in a country which uses A4, then you can rebuild the book. The great thing about the GNU build system is that the book source code which is in Texinfo is also distributed with the program source code, enabling you to do such customizations (hacking).

In order to change the paper size, you will need to have GNU Texinfo installed. Open `doc/gnuastro.texi` with any text editor. This is the source file that created this book. In the first few lines you will see this line:

```
@c@afourpaper
```

In Texinfo, a line is commented with `@c`. Therefore, un-comment this line by deleting the first two characters such that it changes to:

```
@afourpaper
```

Save the file and close it. You can now run

```
$ make pdf
```

and the new PDF book will be available in `SRCdir/doc/gnuastro.pdf`. By changing the `pdf` in `$ make pdf` to `ps` or `dvi` you can have the book in those formats. Note that you can do this for any book that is in Texinfo format, they might not have `@afourpaper` line, so you can add it close to the top of the Texinfo source file.

3.3.4 Known issues

Depending on your operating system and the version of the compiler you are using, you might confront some known problems during the configuration (`$ ./configure`), compilation (`$ make`) and tests (`$ make check`). Here, their solutions are discussed.

- `$ ./configure`: *Configure complains about not finding a library even though you have installed it.* The possible solution is based on how you installed the package:
  - From your distribution's package manager. Most probably this is because your distribution has separated the header files of a library from the library parts. Please also install the 'development' packages for those libraries too. Just add a `-dev` or `-devel` to the end of the package name and re-run the package manager. This will not happen if you install the libraries from source. When installed from source, the headers are also installed.
  - From source. Then your linker is not looking where you installed the library. If you followed the instructions in this chapter, all the libraries will be installed in `/usr/local/lib`. So you have to tell your linker to look in this directory. To do so, add `LDFLAGS=-L/usr/local/lib` to the Gnuastro configure script. If you want to use the libraries for your other programming projects, then export this environment variable similar to the case for `LD_LIBRARY_PATH` explained below.
- `$ make`: *Complains about an unknown function on a non-GNU based operating system.* In this case, please run `$ ./configure` with the `--enable-gnulibcheck` option to see if the problem is from the GNU Portability Library (Gnulib) not supporting your system or if there is a problem in Gnuastro, see Section 3.3.1.1 [Gnuastro configure options], page 37. If the problem is not in Gnulib and after all its tests you get the same complaint from `make`, then please contact us at `bug-gnuastro@gnu.org`. The cause is probably that a function that we have used is not supported by your operating system and we didn't included it along with the source tar ball. If the function is available in Gnulib, it can be fixed immediately.
- `$ make`: *Can't find the headers (.h files) of libraries installed from source.* Similar to the case for `LDFLAGS` (above), your compiler is not looking in the right place, add `CPPFLAGS=-I/usr/local/include` to `./configure` to re-configure Gnuastro, then re-run make.

- `$ make check`: *Only one* `.sh` *test passes, all the rest fail.* It is highly likely that your distribution doesn't look into the `/usr/local/lib` directory when searching for shared libraries. To make sure it is added to the list of directories, run the following command and restart your terminal: (you can ignore the \ and extra space if you type it, it is only necessary if you copy and paste). See Section 3.3.1.2 [Installation directory], page 39, for more details.

  ```
  echo 'export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/local/lib' \
       >> ~/.bashrc
  ```

- `$ make check`: *The tests relying on external programs (for example* `fitstopdf.sh` *fail.)* This is probably due to the fact that the version number of the external programs is too old for the tests we have preformed. Please update the program to a more recent version. For example to create a PDF image, you will need GPL Ghostscript, but older versions do not work, we have successfully tested it on version 9.15. Older versions might cause a failure in the test result.
- `$ make pdf`: *The PDF book cannot be made.* To make a PDF book, you need to have the GNU Texinfo program (like any program, the more recent the better). A working TEX program is also necessary, which you can get from Tex Live[21].
- After `make check`: do not copy the programs' executables to another (for example, the installation) directory manually (using `cp`, or `mv` for example). In the default configuration[22], the program binaries need to link with Gnuastro's shared library which is also built and installed with the programs. Therefore, to run successfully before and after installation, linking modifications need to be made by GNU Libtool at installation time. `make install` does this internally, but a simple copy might give linking errors when you run it. If you need to copy the executables, you can do so after installation.

If your problem was not listed above, please file a bug report (Section 1.7 [Report a bug], page 9).

---

[21] `https://www.tug.org/texlive/`

[22] If you configure Gnuastro with the `--disable-shared` option, then the libraries will be statically linked to the programs and this problem won't exist, see Section 10.1.2 [Linking], page 227.

4 Common program behavior

All the programs in Gnuastro share a set of common behavior mainly to do with user interaction to facilitate their usage. The most basic is how you can configure each program to do what you want: define the input, change parameter/option values, or identify the output. All Gnuastro programs can also read your desired configuration from pre-defined or user-specified files so you don't have to specify all the (sometimes numerous) parameters on the command-line each time you run a program. These files define the "default" program behavior in each directory, for each user, or on each system. In other cases, some programs can greatly benefit from the many threads available in modern CPUs, so here we'll also discuss how you can get the most out of your hardware. Among some other issues, we will also discuss how you can get immediate and distraction-free (without taking your hands off the keyboard!) help, or access to this whole book, on the command-line.

4.1 Command-line

All the programs in Gnuastro are customized through the standard GNU style command-line options. Thus, we'll start by defining this general style that is very common in many command-line tools on unix-like operating systems. Finally, the options that are common to all the programs in Gnuastro are discussed.

The command-line text that you type is passed onto the shell (or program managing the command-line) as a string of characters. See the "Invoking ProgramName" sections in this manual for some examples of commands with each program, for example Section 5.3.1 [Invoking Table], page 94. That string is then broken up into separate *tokens* or *words* by any *metacharacters* (like space, tab, `|`, `>` or `;`) that might exist in the text. To learn more, please see the GNU Bash manual, for the complete list of meta-characters and other GNU Bash definitions (GNU Bash is the most common shell program). Its "Shell Operation" section has a short summary of the steps the shell takes before passing the commands to the program you called.

4.1.1 Arguments and options

On the command-line, the first thing you usually enter is the name of the program you want to run. After that, you can specify two types of input: *arguments* and *options*. In the GNU-style, arguments are those tokens that are not preceded by any hyphens (`-`, see Section 4.1.1.1 [Arguments], page 49). Here is one example:

```
$ astcrop --ra=53.162551 --dec=-27.789676 -w10 hubble-udf.fits
```

In this example, the argument is `hubble-udf.fits`. Arguments are most commonly the input file names containing your data. Options start with one or two hyphens, followed by an identifier for the option (the option's name) and its value (see Section 4.1.1.2 [Options], page 50). Through options you tell the program how to interpret the data. In this example, we are running Section 6.1 [Crop], page 97, to crop a region of width 10 arcseconds centered at the given RA and Dec from the input Hubble Ultra-Deep Field (UDF) FITS image. So options come with an identifier (the option name which is separate from their value).

Arguments can be both mandatory and optional and unlike options they don't have any identifiers (or help from you). Hence, their order might also matter (for example in `cp` which is used for copying one file to another location). The outputs of `--usage` and `--help`

shows which arguments are optional and which are mandatory, see Section 4.7.1 [--usage], page 74. As their name suggests, *options* on the command-line can be considered to be optional and most of the time, you don't have to worry about what order you specify them in. When the order does matter, or the option can be invoked multiple times, it is explicitly mentioned in the "Invoking ProgramName" section of each program.

In case your arguments or option values contain any of the shell's meta-characters, you have to quote them. If there is only one such character, you can use a backslash (\) before it. If there are multiple, it might be easier to simply put your whole argument or option value inside of double quotes ("). In such cases, everything inside the double quotes will be seen as one token or word.

For example, let's say you want to specify the header data unit (HDU) of your FITS file using a complex expression like '`3; images(exposure > 100)`'. If you simply add these after the `--hdu` (`-h`) option, the programs in Gnuastro will read the value to the HDU option as '`3`' and run. Then, Bash will attempt to run a separate command '`images(exposure > 100)`' and complain about a syntax error. This is because the semicolon (`;`) is an 'end of command' character in the shell. To solve this problem you can simply put double quotes around the whole string you want to pass to `--hdu` as seen below:

```
$ astcrop --hdu="3; images(exposure > 100)" FITSimage.fits
```

Alternatively you can put a '\' before every metacharacter in this string, but try doing that, and probably you will agree that the double quotes are much more easier, elegant and readable.

4.1.1.1 Arguments

In Gnuastro, arguments are almost exclusively used as the input data file names. Please consult the first few paragraph of the "Invoking ProgramName" section for each program for a description of what it expects as input, how many arguments, or input data, it accepts, or in what order. Everything particular about how a program treats arguments, is explained under the "Invoking ProgramName" section for that program.

Generally, if there is a standard file name extension for a particular format, that filename extension is used to separate the kinds of arguments. The list below shows the data formats that are recognized in Gnuastro's programs based on their file name endings. Any argument that doesn't end with the specified extensions below is considered to be a text file (usually catalogs, see Section 4.5 [Tables], page 66). In some cases, a program can accept specific formats, for example Section 5.2 [ConvertType], page 87, also accepts `.jpg` images.

- `.fits`: The standard file name ending of a FITS image.
- `.fits.Z`: A FITS image compressed with `compress`.
- `.fits.gz`: A FITS image compressed with GNU zip (gzip).
- `.fits.fz`: A FITS image compressed with `fpack`.
- `.imh`: IRAF format image file.

Through out this book and in the command-line outputs, whenever we want to generalize all such astronomical data formats in a text place holder, we will use `ASTRdata`, we will assume that the extension is also part of this name. Any file ending with these names is directly passed on to CFITSIO to read. Therefore you don't necessarily have to have these

files on your computer, they can also be located on an FTP or HTTP server too, see the CFITSIO manual for more information.

CFITSIO has its own error reporting techniques, if your input file(s) cannot be opened, or read, those errors will be printed prior to the final error by Gnuastro.

4.1.1.2 Options

Command-line options allow configuring the behavior of a program in all GNU/Linux applications for each particular execution on a particular input data. A single option can be called in two ways: *long* or *short*. All options in Gnuastro accept the long format which has two hyphens an can have many characters (for example `--hdu`). Short options only have one hyphen (`-`) followed by one character (for example `-h`). You can see some examples in the list of options in Section 4.1.2 [Common options], page 52, or those for each program's "Invoking ProgramName" section. Both formats are shown for those which support both. First the short is shown then the long.

Usually, the short options are for when you are writing on the command-line and want to save keystrokes and time. The long options are good for shell scripts, where you aren't usually rushing. Long options provide a level of documentation, since they are more descriptive and less cryptic. Usually after a few months of not running a program, the short options will be forgotten and reading your previously written script will not be easy.

Some options need to be given a value if they are called and some don't. You can think of the latter type of options as on/off options. These two types of options can be distinguished using the output of the `--help` and `--usage` options, which are common to all GNU software, see Section 4.7 [Getting help], page 74. In Gnuastro we use the following strings to specify when the option needs a value and what format that value should be in. More specific tests will be done in the program and if the values are out of range (for example negative when the program only wants a positive value), an error will be reported.

`INT` The value is read as an integer.

`FLT` The value is read as a float. There are generally two types, depending on the context. If they are for fractions, they will have to be less than or equal to unity.

`STR` The value is read as a string of characters (for example a file name) or other particular settings like a HDU name, see below.

To specify a value in the short format, simply put the value after the option. Note that since the short options are only one character long, you don't have to type anything between the option and its value. For the long option you either need white space or an `=` sign, for example `-h2`, `-h 2`, `--hdu 2` or `--hdu=2` are all equivalent.

The short format of on/off options (those that don't need values) can be concatenated for example these two hypothetical sequences of options are equivalent: `-a -b -c4` and `-abc4`. As an example, consider the following command to run Crop:

```
$ astcrop -Dr3 --wwidth 3 catalog.txt --deccol=4 ASTRdata
```

The `$` is the shell prompt, `astcrop` is the program name. There are two arguments (`catalog.txt` and `ASTRdata`) and four options, two of them given in short format (`-D`, `-r`) and two in long format (`--width` and `--deccol`). Three of them require a value and one (`-D`) is an on/off option.

If an abbreviation is unique between all the options of a program, the long option names can be abbreviated. For example, instead of typing `--printparams`, typing `--print` or maybe even `--pri` will be enough, if there are conflicts, the program will warn you and show you the alternatives. Finally, if you want the argument parser to stop parsing arguments beyond a certain point, you can use two dashes: `--`. No text on the command-line beyond these two dashes will be parsed.

Gnuastro has two types of options with values, those that only take a single value are the most common type. If these options are repeated or called more than once on the command-line, the value of the last time it was called will be assigned to it. This is very useful when you are testing/experimenting. Let's say you want to make a small modification to one option value. You can simply type the option with a new value in the end of the command and see how the script works. If you are satisfied with the change, you can remove the original option for human readability. If the change wasn't satisfactory, you can remove the one you just added and not worry about forgetting the original value. Without this capability, you would have to memorize or save the original value somewhere else, run the command and then change the value again which is not at all convenient and is potentially cause lots of bugs.

On the other hand, some options can be called multiple times in one run of a program and can thus take multiple values (for example see the `--column` option in Section 5.3.1 [Invoking Table], page 94. In these cases, the order of stored values is the same order that you specified on the command-line.

Gnuastro's programs don't keep any internal default values, so some options are mandatory and if they don't have a value, the program will complain and abort. Most programs have many such options and typing them by hand on every call is impractical. To facilitate the user experience, after parsing the command-line, Gnuastro's programs read special configuration files to get the necessary values for the options you haven't identified on the command-line. These configuration files are fully described in Section 4.2 [Configuration files], page 59.

CAUTION: In specifying a file address, if you want to use the shell's tilde expansion (`~`) to specify your home directory, leave at least one space between the option name and your value. For example use `-o ~/test`, `--output ~/test` or `--output= ~/test`. Calling them with `-o~/test` or `--output=~/test` will disable shell expansion.

CAUTION: If you forget to specify a value for an option which requires one, and that option is the last one, Gnuastro will warn you. But if it is in the middle of the command, it will take the text of the next option or argument as the value which can cause undefined behavior.

NOTE: In some contexts Gnuastro's counting starts from 0 and in others 1. You can assume by default that counting starts from 1, if it starts from 0 for a special option, it will be explicitly mentioned.

4.1.2 Common options

To facilitate the job of the users and developers, all the programs in Gnuastro share some basic command-line options for the options that are common to many of the programs. The full list is classified as Section 4.1.2.1 [Input/Output options], page 52, Section 4.1.2.2 [Processing options], page 54, and Section 4.1.2.3 [Operating mode options], page 55. In some programs, some of the options are irrelevant, but still recognized (you won't get an unrecognized option error, but the value isn't used). Unless otherwise mentioned, these options are identical between all programs.

4.1.2.1 Input/Output options

These options are to do with the input and outputs of the various programs.

`-h STR/INT`
`--hdu=STR/INT`

The name or number of the desired Header Data Unit, or HDU, in the FITS image. A FITS file can store multiple HDUs or extensions, each with either an image or a table or nothing at all (only a header). Note that counting of the extensions starts from 0(zero), not 1(one). Counting from 0 is forced on us by CFITSIO which directly reads the value you give with this option (see Section 3.1.1.2 [CFITSIO], page 26). When specifying the name, case is not important so `IMAGE`, `image` or `ImAgE` are equivalent.

CFITSIO has many capabilities to help you find the extension you want, far beyond the simple extension number and name. See CFITSIO manual's "HDU Location Specification" section for a very complete explanation with several examples. A `#` is appended to the string you specify for the HDU[1] and the result is put in square brackets and appended to the FITS file name before calling CFITSIO to read the contents of the HDU for all the programs in Gnuastro.

`-s STR`
`--searchin=STR`

Where to match/search for columns when the column identifier wasn't a number, see Section 4.5.3 [Selecting table columns], page 71. The acceptable values are `name`, `unit`, or `comment`. This option is only relevant for programs that take table columns as input.

`-I`
`--ignorecase`

Ignore case while matching/searching column meta-data (in the field specified by the `--searchin`). The FITS standard suggests to treat the column names as case insensitive, which is strongly recommended here also but is not enforced. This option is only relevant for programs that take table columns as input.

This option is not relevant to Section 10.2 [BuildProgram], page 230, hence in that program the short option `-I` is used for include directories, not to ignore case.

[1] With the `#` character, CFITSIO will only read the desired HDU into your memory, not all the existing HDUs in the fits file.

`-o STR`
`--output=STR`
: The name of the output file or directory. With this option the automatic output names explained in Section 4.8 [Automatic output], page 77, are ignored.

`-T STR`
`--type=STR`
: The data type of the output depending on the program context. This option isn't applicable to some programs like Section 5.1 [Fits], page 80, and will be ignored by them. The different acceptable values to this option are fully described in Section 4.4 [Numeric data types], page 64.

`-D`
`--dontdelete`
: By default, if the output file already exists, Gnuastro's programs will silently delete it and put their own outputs in its place. When this option is activated, if the output file already exists, the programs will not delete it, will warn you, and will abort.

`-K`
`--keepinputdir`
: In automatic output names, don't remove the directory information of the input file names. As explained in Section 4.8 [Automatic output], page 77, if no output name is specified (with `--output`), then the output name will be made in the existing directory based on your input's file name (ignoring the directory of the input). If you call this option, the directory information of the input will be kept and the automatically generated output name will be in the same directory as the input (usually with a suffix added). Note that his is only relevant if you are running the program in a different directory than the input data.

`-t STR`
`--tableformat=STR`
: The output table's type. This option is only relevant when the output is a table and its format cannot be deduced from its filename. For example, if a name ending in `.fits` was given to `--output`, then the program knows you want a FITS table. But there are two types of FITS tables: FITS ASCII, and FITS binary. Thus, with this option, the program is able to identify which type you want. The currently recognized values to this option are:

  `txt`
  : A plain text table with white-space characters between the columns (see Section 4.5.2 [Gnuastro text table format], page 69).

  `fits-ascii`
  : A FITS ASCII table (see Section 4.5.1 [Recognized table formats], page 67).

  `fits-binary`
  : A FITS binary table (see Section 4.5.1 [Recognized table formats], page 67).

4.1.2.2 Processing options

Some processing steps are common to several programs, so they are defined as common options to all programs. Note that this class of common options is thus necessarily less common between all the programs than those described in Section 4.1.2.1 [Input/Output options], page 52, or Section 4.1.2.3 [Operating mode options], page 55, options. Also, if they are irrelevant for a program, these options will not display in the `--help` output of the program.

`-Z INT[,INT[,...]]`
`--tilesize=[,INT[,...]]`

The size of regular tiles for tessellation, see Section 4.6 [Tessellation], page 72. For each dimension an integer length (in units of data-elements or pixels) is necessary. If the input dimensionality is different from the number of values given to this option, the program will stop with an error. Values must be separated by commas (`,`) and can also be fractions (for example `4/2`). If they are fractions, the result must be an integer, otherwise an error will be printed.

`-M INT[,INT[,...]]`
`--numchannels=INT[,INT[,...]]`

The number of channels for larger input tessellation, see Section 4.6 [Tessellation], page 72. The number and types of acceptable values are similar to `--tilesize`. The only difference is that instead of length, the integers values given to this option represent the *number* of channels, not their size.

`-F FLT`
`--remainderfrac=FLT`

The fraction of remainder size along all dimensions to add to the first tile. See Section 4.6 [Tessellation], page 72, for a complete description. This option is only relevant if `--tilesize` is not exactly divisible by the input dataset's size in a dimension. If the remainder size is larger than this fraction (compared to `--tilesize`), then the remainder size will be added with one regular tile size and divided between two tiles at the start and end of the given dimension.

`--workoverch`

Ignore the channel borders for the high-level job of the given application. As a result, while the channel borders are respected in defining the small tiles (such that no tile will cross a channel border), the higher-level program operation will ignore them, see Section 4.6 [Tessellation], page 72.

`--checktiles`

Make a FITS file with the same dimensions as the input but each pixel is replaced with the ID of the tile that it is associated with. Note that the tile IDs start from 0. See Section 4.6 [Tessellation], page 72, for more on Tiling an image in Gnuastro.

`--oneelempertile`

When showing the tile values (for example with `--checktiles`, or when the program's output is tessellated) only use one element for each tile. This can be useful when only the relative values given to each tile compared to the rest are important or need to be checked. Since the tiles usually have a large number

of pixels within them the output will be much smaller, and so easier to read, write, store, or send.

Note that when the full input size in any dimension is not exactly divisible by the given `--tilesize` in that dimension, the edge tile(s) will have different sizes (in units of the input's size), see `--remainderfrac`. But with this option, all displayed values are going to have the (same) size of one data-element. Hence, in such cases, the image proportions are going to be slightly different with this option.

If your input image is not exaclty divisible by the tile size and you want one value per tile for some higher-level processing, all is not lost though. You can see how many pixels were within each tile (for example to weight the values or discard some for later processing) with Gnuastro's Statistics (see Section 7.1 [Statistics], page 148) as shown below. The output FITS file is going to have two extensions, one with the median calculated on each tile and one with the number of elements that each tile covers. You can then use the `where` operator in Section 6.2 [Arithmetic], page 108, to set the values of all tiles that don't have the regular area to a blank value.

```
$ aststatistics --median --number --ontile input.fits    \
                --oneelempertile --output=o.fits
$ REGULAR_AREA=1600    # Check second extension of ‘o.fits’.
$ astarithmetic o.fits o.fits $REGULAR_AREA ne nan where \
                -h1 -h2
```

Note that if `input.fits` also has blank values, then the median on tiles with blank values will also be ignored with the command above (which is desirable).

`--inteponlyblank`
: When values are to be interpolated, only change the values of the blank elements, keep the non-blank elements untouched.

`--interpnumngb=INT`
: The number of nearby non-blank neighbors to use for interpolation.

4.1.2.3 Operating mode options

Another group of options that are common to all the programs in Gnuastro are those to do with the general operation of the programs. The explanation for those that are not only limited to Gnuastro but are common to all GNU programs start with (GNU option).

`--`
: (GNU option) Stop parsing the command-line. This option can be useful in scripts or when using the shell history. Suppose you have a long list of options, and want to see if removing some of them (to read from configuration files, see Section 4.2 [Configuration files], page 59) can give a better result. If the ones you want to remove are the last ones on the command-line, you don't have to delete them, you can just add `--` before them and if you don't get what you want, you can remove the `--` and get the same initial result.

`--usage`
: (GNU option) Only print the options and arguments and abort. This is very useful for when you know the what the options do, and have just forgot their long/short identifiers, see Section 4.7.1 [`--usage`], page 74.

`-?`
`--help` (GNU option) Print all options with an explanation and abort. Adding this option will print all the options in their short and long formats, also displaying which ones need a value if they are called (with an `=` after the long format followed by a string specifying the format, see Section 4.1.1.2 [Options], page 50). A short explanation is also given for what the option is for. The program will quit immediately after the message is printed and will not do any form of processing, see Section 4.7.2 [`--help`], page 75.

`-V`
`--version`
(GNU option) Print a short message, showing the full name, version, copyright information and program authors and abort. On the first line, it will print the official name (not executable name) and version number of the program. Following this is a blank line and a copyright information. The program will not run.

`-q`
`--quiet` Don't report steps. All the programs in Gnuastro that have multiple major steps will report their steps for you to follow while they are operating. If you do not want to see these reports, you can call this option and only error/warning messages will be printed. If the steps are done very fast (depending on the properties of your input) disabling these reports will also decrease running time.

`--cite` Print the BibTEX entry for Gnuastro and the particular program (if that program comes with a separate paper) and abort. Citations are vital for the continued work on Gnuastro. Gnuastro started and is continued based on separate research projects. So if you find any of the tools offered in Gnuastro to be useful in your research, please use the output of this command to cite the program and Gnuastro in your research paper. Thank you.

Gnuastro is still new, there is no separate paper only devoted to Gnuastro yet. Therefore currently the paper to cite for Gnuastro is the paper for NoiseChisel which is the first published paper introducing Gnuastro to the astronomical community. Upon reaching a certain point, a paper completely devoted to Gnuastro will be published, see Section 1.5.1 [GNU Astronomy Utilities 1.0], page 6.

`-P`
`--printparams`
With this option, Gnuastro's programs will read your command-line options and all the configuration files. If there is no problem (like a missing parameter or a value in the wrong format or range) and immediately before actually running, the programs will print the full list of option names, values and descriptions, sorted and grouped by context and abort. They will also report the version number, the date they were configured on your system and the time they were reported.

As an example, you can give your full command-line options and even the input and output file names and finally just add `-P` to check if all the parameters are finely set. If everything is ok, you can just run the same command (easily

retrieved from the shell history, with the top arrow key) and simply remove the last two characters that showed this option.

Since no program will actually start its processing when this option is called, the otherwise mandatory arguments for each program (for example input image or catalog files) are no longer required when you call this option.

`--config=STR`
: Parse `STR` as a configuration file immediately when this option is confronted (see Section 4.2 [Configuration files], page 59). The `--config` option can be called multiple times in one run of any Gnuastro program on the command-line or in the configuration files. In any case, it will be immediately read (before parsing the rest of the options on the command-line, or lines in a configuration file).

  Note that by definition, later options on the command-line still take precedence over those in these in any configuration file, including the file(s) given to this option. Also see `--lastconfig` and `--onlyversion` on how this option can be used for reproducible results.

`-S`
`--setdirconf`
: Update the current directory configuration file for the Gnuastro program and quit. The full set of command-line and configuration file options will be parsed and options with a value will be written in the current directory configuration file for this program (see Section 4.2 [Configuration files], page 59). If the configuration file or its directory doesn't exist, it will be created. If a configuration file exists it will be replaced (after it, and all other configuration files have been read). In any case, the program will not run.

  This is the recommended method[2] to edit/set the configuration file for all future calls to Gnuastro's programs. It will internally check if your values are in the correct range and type and save them according to the configuration file format, see Section 4.2.1 [Configuration file format], page 59. So if there are unreasonable values to some options, the program will notify you and abort before writing the final configuration file.

  When this option is called, the otherwise mandatory arguments, for example input image or catalog file(s), are no longer mandatory (since the program will not run).

`-U`
`--setusrconf`
: Update the user configuration file and quit (see Section 4.2 [Configuration files], page 59). See explanation under `--setdirconf` for more details.

`--lastconfig`
: This is the last configuration file that must be read. When this option is confronted in any stage of reading the options (on the command-line or in a configuration file), no other configuration file will be parsed, see Section 4.2.2

[2] Alternatively, you can use your favorite text editor.

[Configuration file precedence], page 60, and Section 4.2.3 [Current directory and User wide], page 61. Like all on/off options, on the command-line, this option doesn't take any values. But in a configuration file, it takes the values of `0` or `1`, see Section 4.2.1 [Configuration file format], page 59. If it is present in a configuration file with a value of `0`, then all later occurances of this option will be ignored.

`--onlyversion=STR`

Only run the program if Gnuastro's version is exactly equal to `STR` (see Section 1.5 [Version numbering], page 5). Note that it is not compared as a number, but as a string of characters, so `0`, or `0.0` and `0.00` are different. If the running Gnuastro version is different, then thiss option will report an error and abort as soon as it is confronted on the command-line or in a configuration file. If the running Gnuastro version is the same as `STR`, then the program will run as if this option was not called.

This is useful if you want your results to be exactly reproducible and not mistakenly run with an updated/newer or older version of the program. Besides internal algorithmic/behavior changes in programs, the existence of options or their names might change between versions (especially in these earlier versions of Gnuastro).

Hence, when using this option (probably in a script or in a configuration file), be sure to call it before other options. The benefit is that, when the version differs, the other options won't be parsed and you, or your collaborators/users, won't get errors saying an option in your configuration doesn't exist in the running version of the program.

Here is one example of how this option can be used in conjunction with the `--lastconfig` option. Let's assume that you were satisfied with the results of this command: `astnoisechisel image.fits --detquant=0.95` (along with various options set in various configuration files). You can save the state of NoiseChisel and reproduce that exact result on `image.fits` later by following these steps (the the extra spaces, and `\`, are only for easy readability, if you want to try it out, only one space between each token is enough).

```
$ echo "onlyversion X.XX"             > reproducible.conf
$ echo "lastconfig 1"                >> reproducible.conf
$ astnoisechisel image.fits --detquant=0.95 -P       \
                                     >> reproducible.conf
```

`--onlyversion` was available from Gnuastro 0.0, so putting it immediately at the start of a configuration file will ensure that later, you (or others using different version) won't get a non-recognized option error in case an option was added/removed. `--lastconfig` will inform the installed NoiseChisel to not parse any other configuration files. This is done because we don't want the user's user-wide or system wide option values affecting our results. Finally, with the third command, which has a `-P` (short for `--printparams`), NoiseChisel will print all the option values visible to it (in all the configuration files) and the shell will append them to `reproduce.conf`. Hence, you don't have to worry about remembering the (possibly) different options in the different configuration files.

Afterwards, if you run NoiseChisel as shown below (telling it to read this configuration file with the `--config` option). You can be sure that there will either be an error (for version mis-match) or it will produce exactly the same result that you got before.

```
$ astnoisechisel --config=reproducible.conf
```

`--log` Some programs can generate extra information about their outputs in a log file. When this option is called in those programs, the log file will also be printed. If the program doesn't generate a log file, this option is ignored.

`-N INT`
`--numthreads=INT`
Use `INT` CPU threads when running a Gnuastro program (see Section 4.3 [Multi-threaded operations], page 62). Note that multi-threaded programming is only relevant to some programs. In others, this option will be ignored. If this option is not specified on the command-line or any configuration file, the number of threads will be determined by the programs at configuration time.

4.2 Configuration files

Each program needs a certain number of parameters to run. Supplying all the necessary parameters each time you run the program is very frustrating and prone to errors. Therefore all the programs read the values for the necessary options you have not given in the command line from one of several plain text files (which you can view and edit with any text editor). These files are known as configuration files and are usually kept in a directory named `etc/` according to the file system hierarchy standard[3].

The thing to have in mind is that none of the programs in Gnuastro keep any internal default value. All the values must either be stored in one of the configuration files or explicitly called in the command-line. In case the necessary parameters are not given through any of these methods, the program will print a missing option error and abort. The only exception to this is `--numthreads`, whose default value is determined at run-time using the number of threads available to your system, see Section 4.3 [Multi-threaded operations], page 62. Of course, you can still provide a default value for the number of threads at any of the levels below, but if you don't, the program will not abort. Also note that through automatic output name generation, the value to the `--output` option is also not mandatory on the command-line or in the configuration files for all programs which don't rely on that value as an input[4], see Section 4.8 [Automatic output], page 77.

4.2.1 Configuration file format

The configuration files for each program have the standard program executable name with a '`.conf`' suffix. When you download the source code, you can find them in the same directory as the source code of each program, see Section 11.4 [Program source], page 326.

Any line in the configuration file whose first non-white character is a `#` is considered to be a comment and is ignored. An empty line is also similarly ignored. The long name of the

---

[3] `http://en.wikipedia.org/wiki/Filesystem_Hierarchy_Standard`

[4] One example of a program which uses the value given to `--output` as an input is ConvertType, this value specifies the type of the output through the value to `--output`, see Section 5.2.3 [Invoking ConvertType], page 90.

option should be used as an identifier. The parameter name and parameter value have to be separated by any number of 'white-space' characters: space, tab or vertical tab. By default several space characters are used. If the value of an option has space characters (most commonly for the `hdu` option), then the full value can be enclosed in double quotation signs (`"`, similar to the example in Section 4.1.1 [Arguments and options], page 48). If it is an option without a value in the `--help` output (on/off option, see Section 4.1.1.2 [Options], page 50), then the value should be `1` if it is to be 'on' and `0` otherwise.

In each non-commented and non-blank line, any text after the first two words (option identifier and value) is ignored. If an option identifier is not recognized in the configuration file, the name of the file, the line number of the unrecognized option, and the unrecognized identifier name will be reported and the program will abort. If a parameter is repeated more more than once in the configuration files, accepts only one value, and is not set on the command-line, then only the first value will be used, the rest will be ignored.

You can build or edit any of the directories and the configuration files yourself using any text editor. However, it is recommended to use the `--setdirconf` and `--setusrconf` options to set default values for the current directory or this user, see Section 4.1.2.3 [Operating mode options], page 55. With these options, the values you give will be checked before writing in the configuration file. They will also print a set of commented lines guiding the reader and will also classify the options based on their context and write them in their logical order to be more understandable.

4.2.2 Configuration file precedence

The option values in all the programs of Gnuastro will be filled in the following order. If an option only takes one value which is given in an earlier step, any value for that option in a later step will be ignored. Note that if the `lastconfig` option is specified in any step below, all later files will be ignored (see Section 4.1.2.3 [Operating mode options], page 55). The basic idea behind setting this progressive state of checking for parameter values is that separate users of a computer or separate folders in a user's file system might need different values for some parameters.

In each step, there can also be a configuration file containing the common options in all the programs: `gnuastro.conf` (see Section 4.1.2 [Common options], page 52). If options specific to one program are specified in this file, there will be un-recognized option errors, or unexpected behavior if the option has different behavior in another program. On the other hand, there is no problem with `astprogname.conf` containing common options[5].

1. Command-line options, for a particular run of ProgramName.
2. `.gnuastro/astprogname.conf` is parsed by ProgramName in the current directory.
3. `.gnuastro/gnuastro.conf` is parsed by all Gnuastro programs in the current directory.
4. `$HOME/.local/etc/astprogname.conf` is parsed by ProgramName in the user's home directory (see Section 4.2.3 [Current directory and User wide], page 61).
5. `$HOME/.local/etc/gnuastro.conf` is parsed by all Gnuastro programs in the user's home directory (see Section 4.2.3 [Current directory and User wide], page 61).
6. `prefix/etc/astprogname.conf` is parsed by ProgramName in the system-wide installation directory (see Section 4.2.4 [System wide], page 61, for `prefix`).

[5] As an example, the `--setdirconf` and `--setusrconf` options will also write the common options they have read in their produced `astprogname.conf`.

7. `prefix/etc/gnuastro.conf` is parsed by all Gnuastro programs in the system-wide installation directory (see Section 4.2.4 [System wide], page 61, for `prefix`).

> **Manipulating the order:** You can manipulate this order or add new files with the following two options which are fully described in Section 4.1.2.3 [Operating mode options], page 55:
>
> `--config` Allows you to define any file to be parsed as a configuration file on the command-line or within the any other configuration file. Recall that the file given to `--config` is parsed immediately when this option is confronted (on the command-line or in a configuration file).
>
> `--lastconfig`
> Allows you to stop the parsing of subsequent configuration files. Note that if this option is given in a configuration file, it will be fully read, so its position in the configuration doesn't matter (unlike `--config`).

One example of benefiting from these configuration files can be this: raw telescope images usually have their main image extension in the second FITS extension, while processed FITS images usually only have one extension. If your system-wide default input extension is 0 (the first), then when you want to work with the former group of data you have to explicitly mention it to the programs every time. With this progressive state of default values to check, you can set different default values for the different directories that you would like to run Gnuastro in for your different purposes, so you won't have to worry about this issue any more.

The same can be said about the `gnuastro.conf` files: by specifying a behavior in this single file, all Gnuastro programs in the respective directory, user, or system-wide steps will behave similarly. For example to keep the input's directory when no specific output is given (see Section 4.8 [Automatic output], page 77), or to not delete an existing file if it has the same name as a given output (see Section 4.1.2.1 [Input/Output options], page 52).

4.2.3 Current directory and User wide

For the current (local) and user-wide directories, the configuration files are stored in the hidden sub-directories named `.gnuastro/` and `$HOME/.local/etc/` respectively. Unless you have changed it, the `$HOME` environment variable should point to your home directory. You can check it by running `$ echo $HOME`. Each time you run any of the programs in Gnuastro, this environment variable is read and placed in the above address. So if you suddenly see that your home configuration files are not being read, probably you (or some other program) has changed the value of this environment variable.

Although it might cause confusions like above, this dependence on the `HOME` environment variable enables you to temporarily use a different directory as your home directory. This can come in handy in complicated situations. To set the user or current directory configuration files based on your command-line input, you can use the `--setdirconf` or `--setusrconf`, see Section 4.1.2.3 [Operating mode options], page 55.

4.2.4 System wide

When Gnuastro is installed, the configuration files that are shipped with the distribution are copied into the (possibly system wide) `prefix/etc/` directory. For more details on

prefix, see Section 3.3.1.2 [Installation directory], page 39, (by default it is: /usr/local). This directory is the final place (with the lowest priority) that the programs in Gnuastro will check to retrieve parameter values.

If you remove an option and its value from the system wide configuration files, you either have to specify it in more immediate configuration files or set it each time in the command-line. Recall that none of the programs in Gnuastro keep any internal default values and will abort if they don't find a value for the necessary parameters (except the number of threads and output file name). So even though you might never expect to use an optional option, it safe to have it available in this system-wide configuration file even if you don't intend to use it frequently.

Note that in case you install Gnuastro from your distribution's repositories, prefix will either be set to / (the root directory) or /usr, so you can find the system wide configuration variables in /etc/ or /usr/etc/. The prefix of /usr/local/ is conventionally used for programs you install from source by your self as in Section 1.1 [Quick start], page 1.

4.3 Multi-threaded operations

Some of the programs benefit significantly when you use all the threads your computer's CPU has to offer to your operating system. The number of threads available can be larger than the number of physical (hardware) cores in the CPU (also known as Simultaneous multithreading). For example, in Intel's CPUs (those that implement its Hyper-threading technology) the number of threads is usually double the number of physical cores in your CPU. On a GNU/Linux system, the number of threads available can be found with the command $ nproc command (part of GNU Coreutils).

Gnuastro's programs can find the number of threads available to your system internally at run-time (when you execute the program). However, if a value is given to the --numthreads option, the given number will be used, see Section 4.1.2.3 [Operating mode options], page 55, and Section 4.2 [Configuration files], page 59, for ways to use this option. Thus --numthreads is the only common option in Gnuastro's programs with a value that doesn't have to be specified anywhere on the command-line or in the configuration files.

4.3.1 A note on threads

Spinning off threads is not necessarily the most efficient way to run an application. Creating a new thread isn't a cheap operation for the operating system. It is most useful when the input data are fixed and you want the same operation to be done on parts of it. For example one input image to Crop and multiple crops from various parts of it. In this fashion, the image is loaded into memory once, all the crops are divided between the number of threads internally and each thread cuts out those parts which are assigned to it from the same image. On the other hand, if you have multiple images and you want to crop the same region(s) out of all of them, it is much more efficient to set --numthreads=1 (so no threads spin off) and run Crop multiple times simultaneously, see Section 4.3.2 [How to run simultaneous operations], page 63.

You can check the boost in speed by first running a program on one of the data sets with the maximum number of threads and another time (with everything else the same) and only using one thread. You will notice that the wall-clock time (reported by most programs at their end) in the former is longer than the latter divided by number of physical CPU cores

(not threads) available to your operating system. Asymptotically these two times can be equal (most of the time they aren't). So limiting the programs to use only one thread and running them independently on the number of available threads will be more efficient.

Note that the operating system keeps a cache of recently processed data, so usually, the second time you process an identical data set (independent of the number of threads used), you will get faster results. In order to make an unbiased comparison, you have to first clean the system's cache with the following command between the two runs.

```
$ sync; echo 3 | sudo tee /proc/sys/vm/drop_caches
```

SUMMARY: Should I use multiple threads? Depends:

- If you only have **one** data set (image in most cases!), then yes, the more threads you use (with a maximum of the number of threads available to your OS) the faster you will get your results.
- If you want to run the same operation on **multiple** data sets, it is best to set the number of threads to 1 and use Make, or GNU Parallel, as explained in Section 4.3.2 [How to run simultaneous operations], page 63.

4.3.2 How to run simultaneous operations

There are two[6] approaches to simultaneously execute a program: using GNU Parallel or Make (GNU Make is the most common implementation). The first is very useful when you only want to do one job multiple times and want to get back to your work without actually keeping the command you ran. The second is usually for more important operations, with lots of dependencies between the different products (for example a full scientific research).

GNU Parallel

When you only want to run multiple instances of a command on different threads and get on with the rest of your work, the best method is to use GNU parallel. Surprisingly GNU Parallel is one of the few GNU packages that has no Info documentation but only a Man page, see Section 4.7.4 [Info], page 76. So to see the documentation after installing it please run

```
$ man parallel
```

As an example, let's assume we want to crop a region fixed on the pixels (500, 600) with the default width from all the FITS images in the `./data` directory ending with `sci.fits` to the current directory. To do this, you can run:

```
$ parallel astcrop --numthreads=1 --xc=500 --yc=600 ::: \
  ./data/*sci.fits
```

GNU Parallel can help in many more conditions, this is one of the simplest, see the man page for lots of other examples. For absolute beginners: the backslash (\) is only a line breaker to fit nicely in the page. If you type the whole command in one line, you should remove it.

[6] A third way would be to open multiple terminal emulator windows in your GUI, type the commands separately on each and press `Enter` once on each terminal, but this is far too frustrating, tedious and prone to errors. It's therefore not a realistic solution when tens, hundreds or thousands of operations (your research targets, multiplied by the operations you do on each) are to be done.

Make Make is a program for building "targets" (e.g., files) using "recipes" (a set of operations) when their known "prerequisites" (other files) have been updated. It elegantly allows you to define dependency structures for building your final output and updating it efficiently when the inputs change. It is the most common infra-structure to build software today.

Scientific research methodology is very similar to software developement: you start by testing a hypothesis on a small sample of objects/targets with a simple set of steps. As you are able to get promising results, you improve the method and use it on a larger, more general, sample. In the process, you will confront many issues that have to be corrected (bugs in software development jargon). Make a wonderful tool to manage this style of development. It has been used to make reproducible papers, for example see the reproduction pipeline (`https://gitlab.com/makhlaghi/NoiseChisel-paper`) of the paper introducing Section 7.2 [NoiseChisel], page 163, (one of Gnuastro's programs).

GNU Make[7] is the most common implementation which (similar to nearly all GNU programs, comes with a wonderful manual[8]). Make is very basic and simple, and thus the manual is short (the most important parts are in the first roughly 100 pages) and easy to read/understand.

Make comes with a `--jobs` (`-j`) option which allows you to specify the maximum number of jobs that can be done simultaneously. For example if you have 8 threads available to your operating system. You can run:

```
$ make -j8
```

With this command, Make will process your `Makefile` and create all the targets (can be thousands of FITS images for example) simultaneously on 8 threads, while fully respecting their dependencies (only building a file/target when its prerequisites are successfully built). Make is thus strongly recommended for managing scientific research where robustness, archivability, reproducibility and speed[9] are very important.

4.4 Numeric data types

At the lowest level, the computer stores everything in terms of `1` or `0`. For example, each program in Gnuastro, or each astronomical image you take with the telescope is actually a string of millions of these zeros and ones. The space required to keep a zero or one is the smallest unit of storage, and is known as a *bit*. However, nderstanding and manipulating this string of bits is extremely hard for most people. Therefore, we define packages of these bits along with a standard on how to interpret the bits in each package as a *type*.

The most basic standard for reading the bits is integer numbers ($..., -2, -1, 0, 1, 2, ...$, more bits will give larger limits). The common integer types are 8, 16, 32, and 64 bits

[7] `https://www.gnu.org/software/make/`

[8] `https://www.gnu.org/software/make/manual/`

[9] Besides its multi-threaded capabilities, Make will only re-build those targets that depend on a change you have made, not the whole work. For example, if you have set the prerequisites properly, you can easily test the changing of a parameter on your paper's results without having to re-do everything (which is much faster). This allows you to be much more productive in easily checking various ideas/assumptions of the different stages of your research and thus produce a more robust result for your exciting science.

wide. For each width, there are two standards for reading the bits: signed and unsigned integers. In the former, negative numbers are allowed and in the latter, they aren't. The `unsigned` types thus have larger positive limits (one extra bit), but no negative value. When the context of your work doesn't involve negative numbers (for example counting, where negative is not defined), it is best to use the `unsigned` types. For full numerical range of all integer types, see below.

Another standard of converting a given number of bits to numbers is the floating point standard, this standard can approximately store any real number with a given precision. There are two common floating point types: 32-bit and 64-bit, for single and double precision floating point numbers respectively. The former is sufficient for data with less than 8 significant decimal digits (most astronomical data), while the latter is good for less than 16 significant decimal digits. The representation of real numbers as bits is much more complex than integers. If you are interested, you can start with the Wikipedia article (`https://en.wikipedia.org/wiki/Floating_point`).

With the conversion operators in Gnuastro's Arithmetic, you can change the types of data to each other, which is necessary in some contexts. For example the program/library, that you intend to feed the data into, only accepts floating point values, but you have an integer image. Another situation that conversion can be helpful is when you know that your data only has values that fit within `int8` or `uint16`. However it is currently formatted in the `float64` type. Operations involving floating point or larger integer types are significantly slower than integer or smaller-width types respectively. In the latter case, it also requires much more (by 8 or 4 times in the example above) storage space. So when you confront such situations and want to store/archive/transfter the data, it is best convert them to the most efficient type.

The short and long names for the recognized numeric data types in Gnuastro are listed below. Both short and long names can be used when you want to specify a type. For example, as a value to the common option `--type` (see Section 4.1.2.1 [Input/Output options], page 52), or in the information comment lines of Section 4.5.2 [Gnuastro text table format], page 69. The ranges listed below are inclusive.

`u8`
`uint8`
: 8-bit un-signed integers, range:
$[0 \text{ to } 2^8 - 1]$ or $[0 \text{ to } 255]$.

`i8`
`int8`
: 8-bit signed integers, range:
$[-2^7 \text{ to } 2^7 - 1]$ or $[-127 \text{ to } 127]$.

`u16`
`uint16`
: 16-bit un-signed integers, range:
$[0 \text{ to } 2^{16} - 1]$ or $[0 \text{ to } 65535]$.

`i16`
`int16`
: 16-bit signed integers, range:
$[-2^{15} \text{ to } 2^{15} - 1]$ or $[-32768 \text{ to } 32768]$.

`u32`
`uint32`
: 32-bit un-signed integers, range:
$[0 \text{ to } 2^{32} - 1]$ or $[0 \text{ to } 4294967295]$.

`i32`
`int32` 32-bit signed integers, range:
$[-2^{31}$ to $2^{31} - 1]$ or $[-2147483648$ to $2147483647]$.

`u64`
`uint64` 64-bit un-signed integers, range
$[0$ to $2^{64} - 1]$ or $[0$ to $18446744073709551615]$.

`i64`
`int64` 64-bit signed integers, range:
$[-2^{63}$ to $2^{63} - 1]$ or $[-9223372036854775808$ to $9223372036854775807]$.

`f32`
`float32` 32-bit (single-precision) floating point types. The maximum (minimum is its negative) possible value is 3.402823×10^{38}. Single-precision floating points can accurately represent a floating point number upto ~ 7.2 significant decimals. Given the heavy noise in astronomical data, this is usually more than sufficient for storing results.

`f64`
`float64` 64-bit (double-precision) floating point types. The maximum (minimum is its negative) possible value is $\sim 10^{308}$. Double-precision floating points can accurately represent a floating point number ~ 15.9 significant decimals. This is usually good for processing (mixing) the data internally, for example a sum of single precision data (and later storing the result as `float32`).

Some file formats don't recognize all types. Some file formats don't recognize all the types, for example the FITS standard (see Section 5.1 [Fits], page 80) does not define `uint64` in binary tables or images. When a type is not acceptable for output into a given file format, the respective Gnuastro program or library will let you know and abort. On the command-line, you can use the Section 6.2 [Arithmetic], page 108, program to convert the numerical type of a dataset, in the libraries, you can call `gal_data_copy_to_new_type`.

4.5 Tables

"A table is a collection of related data held in a structured format within a database. It consists of columns, and rows." (from Wikipedia). Each column in the table contains the values of one property and each row is a collection of properties (columns) for one target object. For example, let's assume you have just ran MakeCatalog (see Section 7.3 [MakeCatalog], page 175) on an image to measure some properties for the labeled regions (which might be detected galaxies for example) in the image. For each labeled region (detected galaxy), there will be a *row* which groups its measured properties as *columns*, one column for each property. One such property can be the object's magnitude, which is the sum of pixels with that label, or its center can be defined as the light-weighted average value of those pixels. Many such properties can be derived from the raw pixel values and their position, see Section 7.3.5 [Invoking MakeCatalog], page 185, for a long list.

As a summary, for each labeled region (or, galaxy) we have one *row* and for each measured property we have one *column*. This high-level structure is usually the first step for

higher-level analysis, for example finding the stellar mass or photometric redshift from magnitudes in multiple colors. Thus, tables are not just outputs of programs, infact it is much more common for tables to be inputs of programs. For example, to make a mock galaxy image, you need to feed in the properties of each galaxy into Section 8.1 [MakeProfiles], page 195, for it do the inverse of the process above and make a simulated image from a catalog, see Section 2.2 [Sufi simulates a detection], page 17. In other cases, you can feed a table into Section 6.1 [Crop], page 97, and it will crop out regions centered on the positions within the table, see Section 2.1 [Hubble visually checks and classifies his catalog], page 14. So to end this relatively long introduction, tables play a very important role in astronomy, or generally all branches of data analysis.

In Section 4.5.1 [Recognized table formats], page 67, the currently recognized table formats in Gnuastro are discussed. You can use any of these tables as input or ask for them to be built as output. The most common type of table format is a simple plain text file with each row on one line and columns separated by white space characters, this format is easy to read/write by eye/hand. To give it the full functionality of more specific table types like the FITS tables, Gnuastro has a special convention which you can use to give each column a name, type, unit, and comments, while still being readable by other plain text table readers. This convention is described in Section 4.5.2 [Gnuastro text table format], page 69.

When tables are input to a program, the program reading it needs to know which column(s) it should use for its desired purposes. Gnuastro's programs all follow a similar convention, on the way you can select columns in a table. They are thoroughly discussed in Section 4.5.3 [Selecting table columns], page 71.

4.5.1 Recognized table formats

The list of table formats that Gnuastro can currently read from and write to are decribed below. Each has their own advantage and disadvantages, so a short review of the format is also provided to help you make the best choice based on how you want to define your input tables or later use your output tables.

Plain text table
: This is the most basic and simplest way to create, view, or edit the table by hand on a text editor. The other formats described below are less eye-friendly and have a more formal structure (for easier computer readability). It is fully described in Section 4.5.2 [Gnuastro text table format], page 69.

FITS ASCII tables
: The FITS ASCII table extension is fully in ASCII encoding and thus easily readable on any text editor (assuming it is the only extension in the FITS file). If the FITS file also contains binary extensions (for example an image or binary table extensions), then there will be many hard to print characters. The FITS ASCII format doesn't have new line characters to separate rows. In the FITS ASCII table standard, each row is defined as a fixed number of characters (value to the `NAXIS1` keyword), so to visually inspect it properly, you would have to adjust your text editor's width to this value. All columns start at given character positions and have a fixed width (number of characters).

Numbers in a FITS ASCII table are printed into ASCII format, they are not in binary (that the CPU uses). Hence, they can take a larger space in memory, loose their precision, and take longer to read into memory. If you are dealing with integer type columns (see Section 4.4 [Numeric data types], page 64), another issue with FITS ASCII tables is that the type information for the column will be lost (there is only one integer type in FITS ASCII tables). One problem with the binary format on the other hand is that it isn't portable (different CPUs/compilers) have different standards for translating the zeros and ones. But since ASCII characters are defined on a byte and are well recognized, they are better for portability on those various systems. Gnuastro's plain text table format described below is much more portable and easier to read/write/interpret by humans manually.

Generally, as the name implies, this format is useful for when your table mainly contains ASCII columns (for example file names, or descriptions). They can be useful when you need to include columns with structured ASCII information along with other extensions in one FITS file. In such cases, you can also consider header keywords (see Section 5.1 [Fits], page 80).

FITS binary tables

The FITS binary table is the FITS standard's solution to the issues discussed with keeping numbers in ASCII format as described under the FITS ASCII table title above. Only columns defined as a string type (a string of ASCII characters) are readable in a text editor. The protability problem with binary formats discussed above is mostly solved thanks to the portability of CFITSIO (see Section 3.1.1.2 [CFITSIO], page 26) and the very long history of the FITS format which has been widely used since the 1970s.

In the case of most numbers, storing them in binary format is more memory efficient than ASCII format. For example, to store `-25.72034` in ASCII format, you need 9 bytes/characters. But if you keep this same number (to the approximate precision possible) as a 4-byte (32-bit) floating point number, you can keep/transmit it with less than half the amount of memory. When catalogs contain thousands/millions of rows in tens/hundreds of columns, this can lead to significant improvements in memory/band-width usage. Moreover, since the CPU does its operations in the binary formats, reading the table in and writing it out is also much faster than an ASCII table.

When you are dealing with integer numbers, the compression ratio can be even better, for example if you know all of the values in a column are positive and less than `255`, you can use the `unsigned char` type which only takes one byte! If they are between `-128` and `127`, then you can use the (signed) `char` type. So if you are thoughtful about the limits of your integer columns, you can greatly reduce the size of your file and also the speed at which it is read/written. This can be very useful when sharing your results with collaborators or publishing them. To decrease the file size even more you can name your output as ending in `.fits.gz` so it is also compressed after creation. Just note that compression/decompressing is CPU intensive and can slow down the writing/reading of the file.

Fortunately the FITS Binary table format also accepts ASCII strings as column types (along with the various numerical types). So your dataset can also contain non-numerical columns.

4.5.2 Gnuastro text table format

Plain text files are most generic, portable, and easiest way to (manually) create, (visually) inspect, or (manually) edit a table. In this format, the ending of a row is defined by the new-line character (a line on a text editor). So when you view it on a text editor, every row will occupy one line. The delimiters (or characters separating the columns) are white space characters (space, horizontal tab, vertical tab) and a comma (`,`). The only further requirement is that all rows/lines must have the same number of columns.

The columns don't have to be exactly under each other and the rows can be arbitrarily long with different lengths. For example the following contents in a file would be interpretted as a table with 4 columns and 2 rows, with each element interpretted as a `double` type (see Section 4.4 [Numeric data types], page 64).

```
1     2.234948   128   39.8923e8
2 , 4.454        792     72.98348e7
```

However, the example above has no other information about the columns (its raw data, with no meta-data). To use this table, you have to remember what each column was. Also, when you want to select columns, you have to count their position within the table. This can become frustrating and prone to bad errors (getting the columns wrong) especially as the number of columns increase. It is also bad for sending to a colleague, because they will find it hard to remember/use the columns properly.

To solve these problems in Gnuastro's programs/libraries you aren't limited to using the column's number, see Section 4.5.3 [Selecting table columns], page 71. If the columns have names, units, or comments you can also select your columns based on searches/matches in these fields, for example see Section 5.3 [Table], page 94. Also, in this manner, you can't guide the program reading the table on how to read the numbers. As an example, the first and third columns above can be read as integer types: the first column might be an ID and the third can be the number of pixels an object occupies in an image. So there is no need to read these to columns as a `double` type (which takes more memory, and is slower).

In the bare-minimum example above, you also can't use strings of characters, for example the names of filters, or some other identifier that includes non-numerical characters. In the absence of any information, only numbers can be read robustly. Assuming we read columns with non-numerical characters as string, there would still be the problem that the strings might contain space (or any delimiter) character for some rows. So, each 'word' in the string will be interpretted as a column and the program will abort with an error that the rows don't have the same number of columns.

To correct for these limitations, Gnuastro defines the following convention for storing the table meta-data along with the raw data in one plain text file. The format is primarily designed for ease of reading/writing by eye/fingers, but is also structured enough to be read by a program.

When the first non-white character in a line is `#`, or there are no non-white characters in it, then the line will not be considered as a row of data in the table (this is a pretty standard convention in many programs, and higher level languages). In the former case,

the line is interpretted as a *comment*. If the comment line starts with '# Column N:', then it is assumed to contain information about column N (a number, counting from 1). Comment lines that don't start with this pattern are ignored and you can use them to include any further information you want to store with the table in the text file. A column information comment is assumed to have the following format:

```
# Column N: NAME [UNIT, TYPE, BLANK] COMMENT
```

Any sequence of characters between ':' and '[' will be interpretted as the column name (so it can contain anything except the '[' character). Anything between the ']' and the end of the line is defined as a comment. Within the brackets, anything before the first ',' is the units (physical units, for example km/s, or erg/s), anything before the second ',' is the short type identifier (see below, and Section 4.4 [Numeric data types], page 64). Finally (still within the brackets), any non-white characters after the second ',' are interpretted as the blank value for that column (see Section 6.1.3 [Blank pixels], page 100). Note that blank values will be stored in the same type as the column, not as a string[10].

When a formatting problem occurs (for example you have specified the wrong type code, see below), or the the column was already given meta-data in a previous comment, or the column number is larger than the actual number of columns in the table (the non-commented or empty lines), then the comment information line will be ignored.

When a comment information line can be used, the leading and trailing white space characters will be stripped from all of the elements. For example in this line:

```
# Column 5:  column name   [km/s,    f,-99] Redshift as speed
```

The NAME field will be 'column name', or TYPE will be 'f'. Note how all the white space characters before and after strings are not used, but those in the middle remained. Also, white space characters aren't mandatory, so in the example above BLANK will be '-99'.

Except for the column number (N), the rest of the fields are optional and the column information comments don't have to be in order. In other words, the information for column $N + m$ ($m > 0$) can be given before column N. Also, you don't have to specify information for all columns. Those columns that don't have this information will be interpretted with the default settings (like the case above: values are double precision floating point, and the column has no name, unit, or comment). So these lines are all acceptable for any table (the first one, with nothing but the column number is redundant):

```
# Column 5:
# Column 1: ID [,i] The Clump ID.
# Column 3: mag_f160w [AB mag, f] Magnitude from the F160W filter
```

The data type of the column should be specified with one of the following values:

- For a numeric column, you can use any of the numeric types (and their recognized identifiers) described in Section 4.4 [Numeric data types], page 64.
- 'strN': for strings. The N value identifies the length of the string (how many characters it has). The start of the string on each row is the first non-delimiter character of the column that has the string type. The next N characters will be interpretted as a string and all leading and trailing white space will be removed.

[10] For floating point types, the nan, or inf strings (both not case-sensitive) refer to IEEE NaN (not a number) and infinity values respectively and will be stored as a floating point, so they are acceptable.

If the next column's characters, are closer than N characters to the start of the string column in that line/row, they will be considered part of the string column. If there is a new-line character before the ending of the space given to the string column (in other words, the string column is the last column), then reading of the string will stop, even if the N characters are not complete yet. See tests/table/table.txt for one example. Therefore, the only time you have to pay attention to the positioning and spaces given to the string column is when it is not the last column in the table.

The only limitation in this format is that trailing and leading white space characters will be removed from the columns that are read. In most cases, this is the desired behavior, but if trailing and leading white-spaces are critically important to your analysis, define your own starting and ending characters and remove them after the table has been read. For example in the sample table below, the two '|' characters (which are arbitrary) will remain in the value of the second column and you can remove them manually later. If only one of the leading or trailing white spaces is important for your work, you can only use one of the '|'s.

```
# Column 1: ID [label, uc]
# Column 2: Notes [no unit, str50]
1    leading and trailing white space is ignored here    2.3442e10
2   |         but they will be preserved here        |   8.2964e11
```

Note that the FITS binary table standard does not define the unsigned int and unsigned long types, so if you want to convert your tables to FITS binary tables, use other types. Also, note that in the FITS ASCII table, there is only one integer type (long). So if you convert a Gnuastro plain text table to a FITS ASCII table with the Section 5.3 [Table], page 94, program, the type information for integers will be lost. Conversely if integer types are important for you, you have to manually set them when reading a FITS ASCII table (for example with the Table program when reading/converting into a file, or with the gnuastro/table.h library functions when reading into memory).

4.5.3 Selecting table columns

At the lowest level, the only defining aspect of a column in a table is its number, or position. But selecting columns purely by number is not very convenient and, especially when the tables are large it can be very frustrating and prone to errors. Hence, table file formats (for example see Section 4.5.1 [Recognized table formats], page 67) have ways to store additional information about the columns (meta-data). Some of the most common pieces of information about each column are its *name*, the *units* of data in the it, and a *comment* for longer/informal description of the column's data.

To facilitate research with Gnuastro, you can select columns by matching, or searching in these three fields, besides the low-level column number. To view the full list of information on the columns in the table, you can use the Table program (see Section 5.3 [Table], page 94) with the command below (replace table-file with the filename of your table, if its FITS, you might also need to specify the HDU/extension which contains the table):

```
$ asttable --information table-file
```

Gnuastro's programs need the columns for different purposes, for example in Crop, you specify the column containing the Right Ascension of the crop centers with the --racol option and the column containing the Declination with --deccol. Thus, there is no unified

common option name to select columns for all programs. However, when the program expects the column for a specific context (like the RA and Dec example above), the option names end in the `col` suffix (for example `--racol` and `--deccol`). These options accept values in integer (column number), or string (metadata match/search) format.

If the value can be parsed as a positive integer, it will be seen as the low-level column number. Note that column counting starts from 1, so if you ask for column 0, the respective program will abort with an error. When the value can't be interpretted as an a integer number, it will be seen as a string of characters which will be used to match/search in the table's meta-data. The meta-data field which the value will be compared with can be selected through the `--searchin` option, see Section 4.1.2.1 [Input/Output options], page 52. `--searchin` can take three values: `name`, `unit`, `comment`. The matching will be done following this convention:

- If the value is enclosed in two slashes (for example `-r/RA_/`, or `--racol=/RA_/`), then it is assumed to be a regular expression with the same convention as GNU AWK. GNU AWK has a very well written chapter (`https://www.gnu.org/software/gawk/manual/html_node/Regexp.html`) describing regular expressions, so we we will not continue discussing it here. Regular expressions are a very powerful tool in matching text and useful in many contexts. We thus strongly encourage reviewing this chapter for greatly improving the quality of your work in many cases, not just for searching column meta-data in Gnuastro.
- When the string isn't inclosed between '`/`'s, any column that exactly matches the given value in the given field will be selected.

Note that in both cases, you can ignore the case of alphabetic characters with the `--ignorecase` option, see Section 4.1.2.1 [Input/Output options], page 52. Also, in both cases, multiple columns may be selected with one call to this function. In this case, the order of the selected columns (with one call) will be the same order as they appear in the table.

4.6 Tessellation

It is sometimes necessary to classify the elements in a dataset (for example pixels in an image) into a grid of individual, non-overlapping tiles. For example when background sky gradients are present in an image, you can define a tile grid over the image. When the tile sizes are set properly, the background's variation over each tile will be negligble, allowing you to measure (and subtract) it. In other cases (for example spatial domain convolution in Gnuastro, see Section 6.3 [Convolve], page 117), it might simply be for speed of processing: each tile can be processed independently on a separate CPU thread. In the arts and mathematics, this process is formally known as tessellation (`https://en.wikipedia.org/wiki/Tessellation`).

The size of the regular tiles (in units of data-elements, or pixels in an image) can be defined with the `--tilesize` option. It takes multiple numbers (separated by a comma) which will be the length along the respective dimension (in Fortran/FITS dimension order). Divisions are also acceptable, but must result in an integer. For example `--tilesize=30,40` can be used for an image (a 2D dataset). The regular tile size along the first FITS axis (horizontal when viewed in SAO ds9) will be 30 pixels and along the second it will be 40 pixels. Ideally, `--tilesize` should be selected such that all tiles in the image have exactly

the same size. In other words, that the dataset length in each dimension is divisible by the tile size in that dimension.

However, this is not always possible: the dataset can be any size and every pixel in it is valuable. In such cases, Gnuastro will look at the significance of the remainder length, if it is not significant (for example one or two pixels), then it will just increase the size of the first tile in the respective dimension and allow the rest of the tiles to have the required size. When the remainder is significant (for example one pixel less than the size along that dimension), the remainder will be added to one regular tile's size and the large tile will be cut in half and put in the two ends of the grid/tessellation. In this way, all the tiles in the central regions of the dataset will have the regular tile sizes and the tiles on the edge will be slightly larger/smaller depending on the remainder significance. The fraction which defines the remainder significance along all dimensions can be set through `--remainderfrac`.

The best tile size is directly related to the spatial properties of the property you want to study (for example, gradient on the image). In practice we assume that the gradient is not present over each tile. So if there is a strong gradient (for example in long wavelength ground based images) or the image is of a crowded area where there isn't too much blank area, you have to choose a smaller tile size. A larger mesh will give more pixels and and so the scatter in the results will be less (better statistics).

For raw image processing, a single tesselation/grid is not sufficient. Raw images are the unprocessed outputs of the camera detectors. Modern detectors usually have multiple readout channels each with its own amplifier. For example the Hubble Space Telescope Advanced Camera for Surveys (ACS) has four amplifiers over its full detector area dividing the square field of view to four smaller squares. Ground based image detectors are not exempt, for example each CCD of Subaru Telescope's Hyper Suprime-Cam camera (which has 104 CCDs) has four amplifiers, but they have the same height of the CCD and divide the width by four parts.

The bias current on each amplifier is different, and initial bias subtraction is not perfect. So even after subtracting the measured bias current, you can usually still identify the boundaries of different amplifiers by eye. See Figure 11(a) in Akhlaghi and Ichikawa (2015) for an example. This results in the final reduced data to have non-uniform amplifier-shaped regions with higher or lower background flux values. Such systematic biases will then propagate to all subsequent measurements we do on the data (for example photometry and subsequent stellar mass and star formation rate measurements in the case of galaxies).

Therefore an accurate analysis requires a two layer tesselation: the top layer contains larger tiles, each covering one amplifier channel. For clarity we'll call these larger tiles "channels". The number of channels along each dimension is defined through the `--numchannels`. Each channel is then covered by its own individual smaller tessellation (with tile sizes determined by the `--tilesize` option). This will allow independent analyis of two adjacent pixels from different channels if necessary. If the image is processed or the detector only has one amplifier, you can set the number of channels in both dimension to 1.

The final tesselation can be inspected on the image with the `--checktiles` option that is available to all programs which use tessellation for localized operations. When this option is called, a FITS file with a `_tiled.fits` suffix will be created along with the outputs, see Section 4.8 [Automatic output], page 77. Each pixel in this image has the number of the tile that covers it. If the number of channels in any dimension are larger than unity, you

will notice that the tile IDs are defined such that the first channels is covered first, then the second and so on. For the full list of processing-related common options (including tessellation options), please see Section 4.1.2.2 [Processing options], page 54.

4.7 Getting help

Probably the first time you read this book, it is either in the PDF or HTML formats. These two formats are very convenient for when you are not actually working, but when you are only reading. Later on, when you start to use the programs and you are deep in the middle of your work, some of the details will inevitably be forgotten. Going to find the PDF file (printed or digital) or the HTML webpage is a major distraction.

GNU software have a very unique set of tools for aiding your memory on the command-line, where you are working, depending how much of it you need to remember. In the past, such command-line help was known as "online" help, because they were literally provided to you 'on' the command 'line'. However, nowadays the word "online" refers to something on the internet, so that term will not be used. With this type of help, you can resume your exciting research without taking your hands off the keyboard.

Another major advantage of such command-line based help routines is that they are installed with the software in your computer, therefore they are always in sync with the executable you are actually running. Three of them are actually part of the executable. You don't have to worry about the version of the book or program. If you rely on external help (a PDF in your personal print or digital archive or HTML from the official webpage) you have to check to see if their versions fit with your installed program.

If you only need to remember the short or long names of the options, `--usage` is advised. If it is what the options do, then `--help` is a great tool. Man pages are also provided for those who are use to this older system of documentation. This full book is also available to you on the command-line in Info format. If none of these seems to resolve the problems, there is a mailing list which enables you to get in touch with experienced Gnuastro users. In the subsections below each of these methods are reviewed.

4.7.1 `--usage`

If you give this option, the program will not run. It will only print a very concise message showing the options and arguments. Everything within square brackets (`[]`) is optional. For example here are the first and last two lines of Crop's `--usage` is shown:

```
$ astcrop --usage
Usage: astcrop [-Do?IPqSVW] [-d INT] [-h INT] [-r INT] [-w INT]
            [-x INT] [-y INT] [-c INT] [-p STR] [-N INT] [--deccol=INT]
            ....
            [--setusrconf] [--usage] [--version] [--wcsmode]
            [ASCIIcatalog] FITSimage(s).fits
```

There are no explanations on the options, just their short and long names shown separately. After the program name, the short format of all the options that don't require a value (on/off options) is displayed. Those that do require a value then follow in separate brackets, each displaying the format of the input they want, see Section 4.1.1.2 [Options], page 50. Since all options are optional, they are shown in square brackets, but arguments can also be optional. For example in this example, a catalog name is optional and is only

required in some modes. This is a standard method of displaying optional arguments for all GNU software.

4.7.2 --help

If the command-line includes this option, the program will not be run. It will print a complete list of all available options along with a short explanation. The options are also grouped by their context. Within each context, the options are sorted alphabetically. Since the options are shown in detail afterwards, the first line of the `--help` output shows the arguments and if they are optional or not, similar to Section 4.7.1 [--usage], page 74.

In the `--help` output of all programs in Gnuastro, the options for each program are classified based on context. The first two contexts are always options to do with the input and output respectively. For example input image extensions or supplementary input files for the inputs. The last class of options is also fixed in all of Gnuastro, it shows operating mode options. Most of these options are already explained in Section 4.1.2.3 [Operating mode options], page 55.

The help message will sometimes be longer than the vertical size of your terminal. If you are using a graphical user interface terminal emulator, you can scroll the terminal with your mouse, but we promised no mice distractions! So here are some suggestions:

- `Shift + PageUP` to scroll up and `Shift + PageDown` to scroll down. For most help output this should be enough. The problem is that it is limited by the number of lines that your terminal keeps in memory and that you can't scroll by lines, only by whole screens.
- Pipe to `less`. A pipe is a form of shell re-direction. The `less` tool in Unix-like systems was made exactly for such outputs of any length. You can pipe (`|`) the output of any program that is longer than the screen to it and then you can scroll through (up and down) with its many tools. For example:

  ```
  $ astnoisechisel --help | less
  ```

  Once you have gone through the text, you can quit `less` by pressing the `q` key.
- Redirect to a file. This is a less convenient way, because you will then have to open the file in a text editor! You can do this with the shell redirection tool (`>`):

  ```
  $ astnoisechisel --help > filename.txt
  ```

In case you have a special keyword you are looking for in the help, you don't have to go through the full list. GNU Grep is made for this job. For example if you only want the list of options whose `--help` output contains the word "axis" in Crop, you can run the following command:

```
$ astcrop --help | grep axis
```

If the output of this option does not fit nicely within the confines of your terminal, GNU does enable you to customize its output through the environment variable `ARGP_HELP_FMT`, you can set various parameters which specify the formatting of the help messages. For example if your terminals are wider than 70 spaces (say 100) and you feel there is too much empty space between the long options and the short explanation, you can change these formats by giving values to this environment variable before running the program with the `--help` output. You can define this environment variable in this manner:

```
$ export ARGP_HELP_FMT=rmargin=100,opt-doc-col=20
```

This will affect all GNU programs using GNU C library's `argp.h` facilities as long as the environment variable is in memory. You can see the full list of these formatting parameters in the "Argp User Customization" part of the GNU C library manual. If you are more comfortable to read the `--help` outputs of all GNU software in your customized format, you can add your customizations (similar to the line above, without the `$` sign) to your `~/.bashrc` file. This is a standard option for all GNU software.

4.7.3 Man pages

Man pages were the Unix method of providing command-line documentation to a program. With GNU Info, see Section 4.7.4 [Info], page 76, the usage of this method of documentation is highly discouraged. This is because Info provides a much more easier to navigate and read environment.

However, some operating systems require a man page for packages that are installed and some people are still used to this method of command line help. So the programs in Gnuastro also have Man pages which are automatically generated from the outputs of `--version` and `--help` using the GNU help2man program. So if you run

```
$ man programname
```

You will be provided with a man page listing the options in the standard manner.

4.7.4 Info

Info is the standard documentation format for all GNU software. It is a very useful command-line document viewing format, fully equipped with links between the various pages and menus and search capabilities. As explained before, the best thing about it is that it is available for you the moment you need to refresh your memory on any command-line tool in the middle of your work without having to take your hands off the keyboard. This complete book is available in Info format and can be accessed from anywhere on the command-line.

To open the Info format of any installed programs or library on your system which has an Info format book, you can simply run the command below (change `executablename` to the executable name of the program or library):

```
$ info executablename
```

In case you are not already familiar with it, run `$ info info`. It does a fantastic job in explaining all its capabilities its self. It is very short and you will become sufficiently fluent in about half an hour. Since all GNU software documentation is also provided in Info, your whole GNU/Linux life will significantly improve.

Once you've become an efficient navigator in Info, you can go to any part of this book or any other GNU software or library manual, no matter how long it is, in a matter of seconds. It also blends nicely with GNU Emacs (a text editor) and you can search manuals while you are writing your document or programs without taking your hands off the keyboard, this is most useful for libraries like the GNU C library. To be able to access all the Info manuals installed in your GNU/Linux within Emacs, type `Ctrl-H + i`.

To see this whole book from the beginning in Info, you can run

```
$ info gnuastro
```

If you run Info with the particular program executable name, for example `astcrop` or `astnoisechisel`:

```
$ info astprogramname
```

you will be taken to the section titled "Invoking ProgramName" which explains the inputs and outputs along with the command-line options for that program. Finally, if you run Info with the official program name, for example Crop or NoiseChisel:

```
$ info ProgramName
```

you will be taken to the top section which introduces the program. Note that in all cases, Info is not case sensitive.

4.7.5 help-gnuastro mailing list

Gnuastro maintains the help-gnuastro mailing list for users to ask any questions related to Gnuastro. The experienced Gnuastro users and some of its developers are subscribed to this mailing list and your email will be sent to them immediately. However, when contacting this mailing list please have in mind that they are possibly very busy and might not be able to answer immediately.

To ask a question from this mailing list, send a mail to `help-gnuastro@gnu.org`. Anyone can view the mailing list archives at `http://lists.gnu.org/archive/html/help-gnuastro/`. It is best that before sending a mail, you search the archives to see if anyone has asked a question similar to yours. If you want to make a suggestion or report a bug, please don't send a mail to this mailing list. We have other mailing lists and tools for those purposes, see Section 1.7 [Report a bug], page 9, or Section 1.8 [Suggest new feature], page 11.

4.8 Automatic output

All the programs in Gnuastro are designed such that specifying an output file or directory (based on the program context) is optional. The outputs of most programs are automatically found based on the input and what the program does. For example when you are converting a FITS image named `FITSimage.fits` to a JPEG image, the JPEG image will be saved in `FITSimage.jpg`.

Another very important part of the automatic output generation is that all the directory information of the input file name is stripped off of it. This feature can be disabled with the `--keepinputdir` option, see Section 4.1.2 [Common options], page 52. It is the default because astronomical data are usually very large and organized specially with special file names. In some cases, the user might not have write permissions in those directories. In fact, even if the data is stored on your own computer, it is advised to only grant write permissions to the super user or root. This way, you won't accidentally delete or modify your valuable data!

Let's assume that we are working on a report and want to process the FITS images from two projects (ABC and DEF), which are stored in the sub-directories named `ABCproject/` and `DEFproject/` of our top data directory (`/mnt/data`). The following shell commands show how one image from the former is first converted to a JPEG image through ConvertType and then the objects from an image in the latter project are detected using NoiseChisel. The text after the `#` sign are comments (not typed!).

```
$ pwd                                              # Current location
```

```
/home/usrname/research/report
$ ls                                        # List directory contents
ABC01.jpg
$ ls /mnt/data/ABCproject                                 # Archive 1
ABC01.fits ABC02.fits ABC03.fits
$ ls /mnt/data/DEFproject                                 # Archive 2
DEF01.fits DEF02.fits DEF03.fits
$ astconvertt /mnt/data/ABCproject/ABC02.fits --output=jpg    # Prog 1
$ ls
ABC01.jpg ABC02.jpg
$ astnoisechisel /mnt/data/DEFproject/DEF01.fits              # Prog 2
$ ls
ABC01.jpg ABC02.jpg DEF01_labeled.fits
```

4.9 Output headers

The output FITS files created by Gnuastro's programs have some or all of the following standard keywords to keep the basic date and version information of Gnuastro, its dependencies and the pipeline that is using Gnuastro.

`DATE` The creation time of the FITS file. This date is written directly by CFITSIO and is in UT format.

`COMMIT` Git's commit description from the running directory of Gnuastro's programs. If the running directory is not version controlled or `libgit2` isn't installed (see Section 3.1.2 [Optional dependencies], page 28) then this keyword will not be present. The printed value is equivalent to the output of the following command:

```
git describe --dirty --always
```

If the running directory contains non-committed work, then the stored value will have a '`-dirty`' suffix. This can be very helpful to let you know that the data is not ready to be shared with collaborators or submitted to a journal. You should only share results that are produced after all your work is committed (safely stored in the version controlled history and thus reproducible).

At first sight, version control appears to be mainly a tool for software developers. However progress in a scientific research is almost identical to progress in software development: first you have a rough idea that starts with handful of easy steps. But as the first results appear to be promising, you will have to extend, or generalize, it to make it more robust and work in all the situations your research covers, not just your first test samples. Slowly you will find wrong assumptions or bad implementations that need to be fixed ('bugs' in software development parlance). Finally, when you submit the research to your collaborators or a journal, many comments and suggestions will come in, and you have to address them.

Software developers have created version control systems precisely for this kind of activity. Each significant moment in the project's history is called a "commit", see Section 3.2.2 [Version controlled source], page 32. A snapshot of the project in each "commit" is safely stored away, so you can revert back to it at

a later time, or check changes/progress. This way, you can be sure that your work is reproducible and track the progress and history. With version control, experimentation in the project's analysis is greatly facilitated, since you can easily revert back if a brainstorm test procedure fails.

One important feature of version control is that the research result (FITS image, table, report or paper) can be stamped with the unique commit information that produced it. This information will enable you to exactly reproduce that same result later, even if you have made changes/progress. For one example of a research paper's reproduction pipeline, please see the reproduction pipeline (`https://gitlab.com/makhlaghi/NoiseChisel-paper`) of the paper (`https://arxiv.org/abs/1505.01664`) describing Section 7.2 [NoiseChisel], page 163.

`CFITSIO` The version of CFITSIO used (see Section 3.1.1.2 [CFITSIO], page 26).

`WCSLIB` The version of WCSLIB used (see Section 3.1.1.3 [WCSLIB], page 28). Note that older versions of WCSLIB do not report the version internally. So this is only available if you are using more recent WCSLIB versions.

`GSL` The version of GNU Scientific Library that was used, see Section 3.1.1.1 [GNU Scientific library], page 26.

`GNUASTRO` The version of Gnuastro used (see Section 1.5 [Version numbering], page 5).

Here is one example of the last few lines of an example output.

```
              / Versions and date
DATE    = '...'                / file creation date
COMMIT  = 'v0-8-g547f6eb'      / Commit description in running dir.
CFITSIO = '3.41    '           / CFITSIO version.
WCSLIB  = '5.16    '           / WCSLIB version.
GSL     = '2.3     '           / GNU Scientific Library version.
GNUASTRO= '0.3'                / GNU Astronomy Utilities version.
END
```

5 Data containers

The most low-level and basic property of a dataset is how it is stored. To process, archive and transmit the data, you need a container to store it first. From the start of the computer age, different formats have been defined to store data, optimized for pariticular applications. One format/container can never be useful for all applications: the storage defines the application and vice-versa. In astronomy, the Flexible Image Transport System (FITS) standard has become the most common format of data storage and transmission. It has many useful features, for example multiple sub-containers (also known as extensions or header data units, HDUs) within one file, or support for tables as well as images. Each HDU can store an independent dataset and its corresponding meta-data. Therefore, Gnuastro has one program (see Section 5.1 [Fits], page 80) specifically designed to manipulate FITS HDUs and the meta-data (header keywords) in each HDU.

Your astronomical research does not just involve data analysis (where the FITS format is very useful). For example you want to demonstrate your raw and processed FITS images or spectra as figures within slides, reports, or papers. The FITS format is not defined for such applications. Thus, Gnuastro also comes with the ConvertType program (see Section 5.2 [ConvertType], page 87) which can be used to convert a FITS image to and from (where possible) other formats like plain text and JPEG (which allow two way conversion), along with EPS and PDF (which can only be created from FITS, not the other way round).

Finally, the FITS format is not just for images, it can also store tables. Binary tables in particular can be very efficient in storing catalogs that have more than a few tens of columns and rows. However, unlike images (where all elements/pixels have one data type), tables contain multiple columns and each column can have different properties: independent data types (see Section 4.4 [Numeric data types], page 64) and meta-data. In practice, each column can be viewed as a separate container that is grouped with others in the table. The only shared property of the colums in a table is thus the number of elements they contain. To allow easy inspection/manipulation of table columns, Gnuastro has the Table program (see Section 5.3 [Table], page 94). It can be used to select certain table columns in a FITS table and see them as a human readable output on the command-line, or to save them into another plain text or FITS table.

5.1 Fits

The "Flexible Image Transport System", or FITS, is by far the most common data container format in astronomy. Although the full name of the standard invokes the idea that it is only for images, it also contains very complete and robust features for tables. It started off in the 1970s and was formally published as a standard in 1981, it was adopted by the International Astronomical Union (IAU) in 1982 and an IAU working group to maintain its future was defined in 1988. The FITS 2.0 and 3.0 standards were approved in 2000 and 2008 respectively, and the 4.0 draft has also been released recently, please see the FITS standard document webpage (`https://fits.gsfc.nasa.gov/fits_standard.html`) for the full text of all versions. Also see the FITS 3.0 standard paper (`https://doi.org/10.1051/0004-6361/201015362`) for a nice introduction and history along with the full standard.

Other formats, for example a JPEG image, only have one image/dataset per file, however one great advantage of the FITS standard is that it allows you to keep multiple datasets

(images or tables along with their meta-data) in one file. Each data + metadata is known as an extension, or more formally a header data unit or HDU, in the FITS standard. In theory the HDUs can be completely independent: you can have multiple images of different dimensions/sizes or tables as separate extensions in one file. However, while the standard doesn't impose any constraints on the relation between the datasets, it is strongly encouraged to group data that are contextually related with each other in one file. For example an image and the table/catalog of objects and their measured properties in that image. Another example can be multiple images of one patch of sky in different colors (filters).

As discussed above, the extensions in a FITS file can be completely independent. To keep some information (meta-data) about the group of extensions in the FITS file, the community has adopted the following convension: put no data in the first extension, so it is just meta-data. This extension can thus be used to store Meta-data regarding the whole file (grouping of extensions). Subsequent extensions may contain data along with their own separate meta-data. All of Gnuastro's programs also follow this convention: the main dataset (image or table) is in the second extension. See the example list of extension properties in Section 5.1.1 [Invoking Fits], page 82.

The meta-data contain information about the data, for example which region of the sky an image corresponds to, the units of the data, what telescope, camera, and filter the data were taken with, the software that produced it, or it observation or processing date. Hence without the meta-data, the raw dataset is practically just a collection of numbers and really hard to understand, or connect with the real world (other datasets). It is thus strongly encouraged to supplement your data (at any level of processing) with as much meta-data about your processing/science as possible.

The meta-data of a FITS file is in ASCII format, which can be easily viewed or edited with a text editor or on the command-line. Each meta-data element (known as a keyword generally) is composed of a name, value, units and comments (the last two are optional). For example below you can see three FITS meta-data keywords for specifying the world coordinate system (WCS, or its location in the sky) of a dataset:

```
LATPOLE =           -27.805089 / [deg] Native latitude of celestial pole
RADESYS = 'FK5'                / Equatorial coordinate system
EQUINOX =               2000.0 / [yr] Equinox of equatorial coordinates
```

However, there are some limitations which discourage viewing/editing the keywords with text editors. For example there is a fixed length of 80 characters for each keyword (its name, value, units and comments) and there are no new-line characters, so on a text editor all the keywords are seen in one line. Also, the meta-data keywords are immediately followed by the data which are commonly in binary format and will show up as strange looking characters on a text editor, and significantly slowing down the processor.

Gnuastro's Fits program was designed to allow easy manipulation of FITS extensions and meta-data keywords on the command-line while conforming fully with the FITS standard. For example you can copy or cut (copy and remove) HDUs/extensions from one FITS file to another, or completely delete them. It also has features to delete, add, or edit meta-data keywords within one HDU.

5.1.1 Invoking Fits

Fits can print or manipulate the FITS file HDUs (extensions), meta-data keywords in a given HDU. The executable name is `astfits` with the following general template

```
$ astfits [OPTION...] ASTRdata
```

One line examples:

```
## View general information about every extension:
$ astfits image.fits

## Print the header keywords in the second HDU (counting from 0):
$ astfits image.fits -h1 --printallkeys

## Only print header keywords that contain 'NAXIS':
$ astfits image.fits -h1 --printallkeys | grep NAXIS

## Copy a HDU from input.fits to out.fits:
$ astfits input.fits --copy=hdu-name --output=out.fits

## Update the OLDKEY keyword value to 153.034:
$ astfits --update=OLDKEY,153.034,"Old keyword comment"

## Delete one COMMENT keyword and add a new one:
$ astfits --delete=COMMENT --comment="Anything you like ;-)."

## Write two new keywords with different values and comments:
$ astfits --write=MYKEY1,20.00,"An example keyword" --write=MYKEY2,fd
```

When no action is requested (and only a file name is given), Fits will print a list of information about the extension(s) in the file. This information includes the HDU number, HDU name (`EXTNAME` keyword), type of data (see Section 4.4 [Numeric data types], page 64, and the number of data elements it contains (size along each dimension for images and table rows and columns). You can use this to get a general idea of the contents of the FITS file and what HDU to use for further processing, either with the Fits program or any other Gnuastro program.

Here is one example of information about a FITS file with four exensions: the first extension has no data, it is a purely meta-data HDU (commonly used to keep meta-data about the whole file, or grouping of extensions, see Section 5.1 [Fits], page 80). The second extension is an image with name `IMAGE` and single precision floating point type (`float32`, see Section 4.4 [Numeric data types], page 64), it has 4287 pixels along its first (horizontal) axis and 4286 pixels along its second (vertical) axis. The third extension is also an image with name `MASK`. It is in 2-byte integer format (`int16`) which is commonly used to keep information about pixels (for example to identify which ones were saturated, or which ones had cosmic rays and so on), note how it has the same size as the `IMAGE` extension. The third extension is a binary table called `CATALOG` which has 12371 rows and 5 columns (it probably contains information about the sources in the image).

```
GNU Astronomy Utilities X.X
Run on Day Month DD HH:MM:SS YYYY
```

```
-----
HDU (extension) information: ‘image.fits’.
 Column 1: Index (counting from 0).
 Column 2: Name (‘EXTNAME’ in FITS standard).
 Column 3: Image data type or ‘table’ format (ASCII or binary).
 Column 4: Size of data in HDU.
-----
0      n/a             uint8           0
1      IMAGE           float32         4287x4286
2      MASK            int16           4287x4286
3      CATALOG         table_binary    12371x5
```

The operating mode and input/output options to Fits are similar to the other programs and fully described in Section 4.1.2 [Common options], page 52. The options particular to Fits can be divided into two groups: 1) those related to modifying HDUs or extensions (see Section 5.1.1.1 [HDU manipulation], page 83), and 2) those related to viewing/modifying meta-data keywords (see Section 5.1.1.2 [Keyword manipulation], page 84). These two classes of options cannot be called together in one run: you can either work on the extensions or meta-data keywords in any instance of Fits.

5.1.1.1 HDU manipulation

Each header data unit, or HDU (also known as an extension), in a FITS file is an independent dataset (data + meta-data). Multiple HDUs can be stored in one FITS file, see Section 5.1 [Fits], page 80. The HDU modifying options to the Fits program are listed below.

These options may be called multiple times in one run. If so, the extensions will be copied from the input FITS file to the output FITS file in the given order (on the command-line and also in configuration files, see Section 4.2.2 [Configuration file precedence], page 60). If the separate classes are called together in one run of Fits, then first `--copy` is run (on all specified HDUs), followed `--cut` (again on all specified HDUs), and then `--remove` (on all specified HDUs).

The `--copy` and `--cut` options need an output FITS file (specified with the `--output` option). If the output file exists, then the specified HDU will be copied following the last extension of the output file (its existing in it HDUs will be untouched). Thus after Fits finishes, the copied HDU will be the last HDU of the output file. If no output file name is given, then automatic output will be used to store the HDUs given to this option (see Section 4.8 [Automatic output], page 77).

`-C STR`
`--copy=STR`
: Copy the specified extension into the output file, see explanations above.

`-k STR`
`--cut=STR`
: Cut (copy to output, remove from input) the specified extension into the output file, see explanations above.

`-R STR`
`--remove=STR`
: Remove the specified HDU from the input file.

5.1.1.2 Keyword manipulation

The meta-data in each header data unit, or HDU (also known as extension, see Section 5.1 [Fits], page 80) is stored as "keyword"s. Each keyword consists of a name, value, unit, and comments. The Fits program (see Section 5.1 [Fits], page 80) options related to viewing and manipulating keywords in a FITS HDU are described below.

To see the full list of keywords in a FITS HDU, you can use the `--printallkeys` option. If any of the keywords are to be modified, the headers of the input file will be changed. If you want to keep the original FITS file or HDU, it is easiest to create a copy first and then run Fits on that. In the FITS standard, keywords are always uppercase. So case does not matter in the input or output keyword names you specify.

Most of the options can accept multiple instances in one command. For example you can add multiple keywords to delete by calling `--delete` multiple times, since repeated keywords are allowed, you can even delete the same keyword multiple times. The action of such options will start from the top most keyword.

The precedence of operations are described below. Note that while the order within each class of actions is preserved, the order of individual actions is not. So irrespective of what order you called `--delete` and `--update`. First, all the delete operations are going to take effect then the update operations.

1. `--delete`
2. `--rename`
3. `--update`
4. `--write`
5. `--asis`
6. `--history`
7. `--comment`
8. `--date`
9. `--printallkeys`

All possible syntax errors will be reported before the keywords are actually written. FITS errors during any of these actions will be reported, but Fits won't stop until all the operations are complete. If `--quitonerror` is called, then Fits will immediately stop upon the first error.

If you want to inspect only a certain set of header keywords, it is easiest to pipe the output of the Fits program to GNU Grep. Grep is a very powerful and advanced tool to search strings which is precisely made for such situations. For example if you only want to check the size of an image FITS HDU, you can run:

```
$ astfits input.fits | grep NAXIS
```

FITS STANDARD KEYWORDS: Some header keywords are necessary for later operations on a FITS file, for example BITPIX or NAXIS, see the FITS standard for their full list. If you modify (for example remove or rename) such keywords, the FITS file extension might not be usable any more. Also be careful for the world coordinate system keywords, if you modify or change their values, any future world coordinate system (like RA and Dec) measurements on the image will also change.

The keyword related options to the Fits program are fully described below.

`-a STR`
`--asis=STR`
Write `STR` exactly into the FITS file header with no modifications. If it does not conform to the FITS standards, then it might cause trouble, so please be very careful with this option. If you want to define the keyword from scratch, it is best to use the `--write` option (see below) and let CFITSIO worry about the standards. The best way to use this option is when you want to add a keyword from one FITS file to another unchanged and untouched. Below is an example of such a case that can be very useful sometimes (on the command-line or in scripts):

```
$ key=$(astfits firstimage.fits | grep KEYWORD)
$ astfits --asis="$key" secondimage.fits
```

In particular note the double quotation signs (`"`) around the reference to the `key` shell variable (`$key`), since FITS keywords usually have lots of space characters, if this variable is not quoted, the shell will only give the first word in the full keyword to this option, which will definitely be a non-standard FITS keyword and will make it hard to work on the file afterwords. See the "Quoting" section of the GNU Bash manual for more information if your keyword has the special characters `$`, `‘`, or `\`.

`-d STR`
`--delete=STR`
Delete one instance of the `STR` keyword from the FITS header. Multiple instances of `--delete` can be given (possibly even for the same keyword, when its repeated in the meta-data). All keywords given will be removed from the headers in the same given order. If the keyword doesn't exist, Fits will give a warning and return with a non-zero value, but will not stop. To stop as soon as an error occurs, run with `--quitonerror`.

`-r STR`
`--rename=STR`
Rename a keyword to a new value. `STR` contains both the existing and new names, which should be separated by either a comma (`,`) or a space character. Note that if you use a space character, you have to put the value to this option within double quotation marks (`"`) so the space character is not interpreted as an option separator. Multiple instances of `--rename` can be given in one command. The keywords will be renamed in the specified order. If the keyword doesn't exist, Fits will give a warning and return with a non-zero value, but will not stop. To stop as soon as an error occurs, run with `--quitonerror`.

`-u STR`
`--update=STR`
Update a keyword, its value, its comments and its units in the format described below. If there are multiple instances of the keyword in the header, they will be changed from top to bottom (with multiple `--update` options).

The format of the values to this option can best be specified with an example:

```
--update=KEYWORD,value,"comments for this keyword",unit
```

If there is a writing error, Fits will give a warning and return with a non-zero value, but will not stop. To stop as soon as an error occurs, run with `--quitonerror`.

The value can be any numerical or string value[1]. Other than the `KEYWORD`, all the other values are optional. To leave a given token empty, follow the preceding comma (`,`) immediately with the next. If any space character is present around the commas, it will be considered part of the respective token. So if more than one token has space characters within it, the safest method to specify a value to this option is to put double quotation marks around each individual token that needs it. Note that without double quotation marks, space characters will be seen as option separators and can lead to undefined behavior.

`-w STR`
`--write=STR`
Write a keyword to the header. For the possible value input formats, comments and units for the keyword, see the `--update` option above.

`-H STR`
`--history STR`
Add a `HISTORY` keyword to the header. The string given to this keyword will be separated into multiple keyword cards if it is longer than 70 characters. With each run only one value for the `--history` option will be read. If there are multiple, it is the last one. If there is an error, Fits will give a warning and return with a non-zero value, but will not stop. To stop as soon as an error occurs, run with `--quitonerror`.

`-c STR`
`--comment STR`
Add a `COMMENT` keyword to the header. Similar to the explanation for `--history` above. If there is a writing error, Fits will give a warning and return with a non-zero value, but will not stop. To stop as soon as an error occurs, run with `--quitonerror`.

`-t`
`--date` Put the current date and time in the header. If the `DATE` keyword already exists in the header, it will be updated. If there is a writing error, Fits will give a warning and return with a non-zero value, but will not stop. To stop as soon as an error occurs, run with `--quitonerror`.

`-p`
`--printall`
Print all the keywords in the specified FITS extension (HDU) on the command-line.

[1] Some tricky situations arise with values like '`87095e5`', if this was intended to be a number it will be kept in the header as `8709500000` and there is no problem. But this can also be a shortened Git commit hash. In the latter case, it should be treated as a string and stored as it is written. Commit hashes are very important in keeping the history of a file during your research and such values might arise without you noticing them in your reproduction pipeline. One solution is to use `git describe` instead of the short hash alone. A less recommended solution is to add a space after the commit hash and Fits will write the value as '`87095e5 `' in the header. If you later compare the strings on the shell, the space character will be ignored by the shell in the latter solution and there will be no problem.

`-Q`
`--quitonerror`
: Quit if any of the operations above are not successful. By default if an error occurs, Fits will warn the user of the faulty keyword and continue with the rest of actions.

5.2 ConvertType

The formats of astronomical data were defined mainly for archiving and processing. In other situations, the data might be useful in other formats. For example, when you are writing a paper or report or if you are making slides for a talk, you can't use a FITS image. Other image formats should be used. In other cases you might want your pixel values in a table format as plain text for input to other programs that don't recognize FITS, or to include as a table in your report. ConvertType is created for such situations. The various types will increase with future updates and based on need.

The conversion is not only one way (from FITS to other formats), but two ways (except the EPS and PDF formats). So you can convert a JPEG image or text file into a FITS image. Basically, other than EPS, you can use any of the recognized formats as different color channel inputs to get any of the recognized outputs. So before explaining the options and arguments, first a short description of the recognized files types will be given followed a short introduction to digital color.

5.2.1 Recognized file formats

The various standards and the file name extensions recognized by ConvertType are listed below.

FITS or IMH
: Astronomical data are commonly stored in the FITS format (and in older data sets in IRAF `.imh` format), a list of file name suffixes which indicate that the file is in this format is given in Section 4.1.1.1 [Arguments], page 49.

: Each extension of a FITS image only has one value per pixel, so when used as input, each input FITS image contributes as one color channel. If you want multiple extensions in one FITS file for different color channels, you have to repeat the file name multiple times and use the `--hdu`, `--hdu2`, `--hdu3` or `--hdu4` options to specify the different extensions.

JPEG
: The JPEG standard was created by the Joint photographic experts group. It is currently one of the most commonly used image formats. Its major advantage is the compression algorithm that is defined by the standard. Like the FITS standard, this is a raster graphics format, which means that it is pixelated.

: A JPEG file can have 1 (for grayscale), 3 (for RGB) and 4 (for CMYK) color channels. If you only want to convert one JPEG image into other formats, there is no problem, however, if you want to use it in combination with other input files, make sure that the final number of color channels does not exceed four. If it does, then ConvertType will abort and notify you.

: The file name endings that are recognized as a JPEG file for input are: `.jpg`, `.JPG`, `.jpeg`, `.JPEG`, `.jpe`, `.jif`, `.jfif` and `.jfi`.

EPS
: The Encapsulated PostScript (EPS) format is essentially a one page PostScript file which has a specified size. PostScript also includes non-image data, for example lines and texts. It is a fully functional programming language to describe a document. Therefore in ConvertType, EPS is only an output format and cannot be used as input. Contrary to the FITS or JPEG formats, PostScript is not a raster format, but is categorized as vector graphics.

  The Portable Document Format (PDF) is currently the most common format for documents. Some believe that PDF has replaced PostScript and that PostScript is now obsolete. This view is wrong, a PostScript file is an actual plain text file that can be edited like any program source with any text editor. To be able to display its programmed content or print, it needs to pass through a processor or compiler. A PDF file can be thought of as the processed output of the compiler on an input PostScript file. PostScript, EPS and PDF were created and are registered by Adobe Systems.

  With these features in mind, you can see that when you are compiling a document with TEX or LATEX, using an EPS file is much more low level than a JPEG and thus you have much greater control and therefore quality. Since it also includes vector graphic lines we also use such lines to make a thin border around the image to make its appearance in the document much better. No matter the resolution of the display or printer, these lines will always be clear and not pixelated. In the future, addition of text might be included (for example labels or object IDs) on the EPS output. However, this can be done better with tools within TEX or LATEX such as PGF/Tikz[2].

  If the final input image (possibly after all operations on the flux explained below) is a binary image or only has two colors of black and white (in segmentation maps for example), then PostScript has another great advantage compared to other formats. It allows for 1 bit pixels (pixels with a value of 0 or 1), this can decrease the output file size by 8 times. So if a grayscale image is binary, ConvertType will exploit this property in the EPS and PDF (see below) outputs.

  The standard formats for an EPS file are `.eps`, `.EPS`, `.epsf` and `.epsi`. The EPS outputs of ConvertType have the `.eps` suffix.

PDF
: As explained above, a PDF document is a static document description format, viewing its result is therefore much faster and more efficient than PostScript. To create a PDF output, ConvertType will make a PostScript page description and convert that to PDF using GPL Ghostscript. The suffixes recognized for a PDF file are: `.pdf`, `.PDF`. If GPL Ghostscript cannot be run on the PostScript file, it will remain and a warning will be printed.

`blank`
: This is not actually a file type! But can be used to fill one color channel with a blank value. If this argument is given for any color channel, that channel will not be used in the output.

---

[2] `http://sourceforge.net/projects/pgf/`

Plain text Plain text files have the advantage that they can be viewed with any text editor or on the command-line. Most programs also support input as plain text files. As input, each plain text file is considered to contain one color channel.

In ConvertType, the recognized extensions for plain text files are `.txt` and `.dat`. As described in Section 5.2.3 [Invoking ConvertType], page 90, if you just give these extensions, (and not a full filename) as output, then automatic output will be preformed to determine the final output name (see Section 4.8 [Automatic output], page 77). Besides these, when the format of a file cannot be recognized from its name, ConvertType will fall back to plain text mode. So you can use any name (even without an extension) for a plain text input or output. Just note that when the suffix is not recognized, automatic output will not be preformed.

The basic input/output on plain text images is very similar to how tables are read/written as discribed in Section 4.5.2 [Gnuastro text table format], page 69. Simply put, the restrictions are very loose, and there is a convention to define a name, units, data type (see Section 4.4 [Numeric data types], page 64), and comments for the data in a commented line. The only difference is that as a table, a text file can contain many datasets (columns), but as a 2D image, it can only contain one dataset. As a result, only one information comment line is necessary for a 2D image, and instead of the starting '`# Column N`' (`N` is the column number), the information line for a 2D image must start with '`# Image 1`'. When ConvertType is asked to output to plain text file, this information comment line is written before the image pixel values.

When converting an image to plain text, consider the fact that if the image is large, the number of columns in each line will become very large, possibly making it very hard to open in some text editors.

Standard output (command-line)
This is very similar to the plain text output, but instead of creating a file to keep the printed values, they are printed on the command line. This can be very useful when you want to redirect the results directly to another program in one command with no intermediate file. The only difference is that only the pixel values are printed (with no information comment line). To print to the standard output, set the output name to '`stdout`'.

5.2.2 Color

An image is a two dimensional array of 2 dimensional elements called pixels. If each pixel only has one value, the image is known as a grayscale image and no color is defined. The range of values in the image can be interpreted as shades of any color, it is customary to use shades of black or grayscale. However, to produce the color spectrum in the digital world, several primary colors must be mixed. Therefore in a color image, each pixel has several values depending on how many primary colors were chosen. For example on the digital monitor or color digital cameras, all colors are built by mixing the three colors of Red-Green-Blue (RGB) with various proportions. However, for printing on paper, standard printers use the Cyan-Magenta-Yellow-Key (CMYK, Key=black) color space. Therefore when printing an RGB image, usually a transformation of color spaces will be necessary.

In a colored digital camera, a color image is produced by dividing the pixel's area between three colors (filters). However in astronomy due to the intrinsic faintness of most of the targets, the collecting area of the pixel is very important for us. Hence the full area of the pixel is used and one value is stored for that pixel in the end. One color filter is used for the whole image. Thus a FITS image is inherently a grayscale image and no color can be defined for it.

One way to represent a grayscale image in different color spaces is to use the same proportions of the primary colors in each pixel. This is the common way most FITS image converters work: they fill all the channels with the same values. The downside is two fold:

- Three (for RGB) or four (for CMYK) values have to be stored for every pixel, this makes the output file very heavy (in terms of bytes).
- If printing, the printing errors of each color channel can make the printed image slightly more blurred than it actually is.

To solve both these problems, the best way is to save the FITS image into the black channel of the CMYK color space. In the RGB color space all three channels have to be used. The JPEG standard is the only common standard that accepts CMYK color space, that is why currently only the JPEG standard is included and not the PNG standard for example.

The JPEG and EPS standards set two sizes for the number of bits in each channel: 8-bit and 12-bit. The former is by far the most common and is what is used in ConvertType. Therefore, each channel should have values between 0 to $2^8 - 1 = 255$. From this we see how each pixel in a grayscale image is one byte (8 bits) long, in an RGB image, it is 3bytes long and in CMYK it is 4bytes long. But thanks to the JPEG compression algorithms, when all the pixels of one channel have the same value, that channel is compressed to one pixel. Therefore a Grayscale image and a CMYK image that has only the K-channel filled are approximately the same file size.

5.2.3 Invoking ConvertType

ConvertType will convert any recognized input file type to any specified output type. The executable name is `astconvertt` with the following general template

```
$ astconvertt [OPTION...] InputFile [InputFile2] ... [InputFile4]
```

One line examples:

```
## Convert an image in FITS to PDF:
$ astconvertt image.fits --output=pdf

## Convert an image in JPEG to FITS:
$ astconvertt image.jpg -ogalaxy.fits

## Use three plain text 2D arrays to create an RGB JPEG output:
$ astconvertt f1.txt f2.txt f3.fits -o.jpg

## Use two images and one blank for an RGB EPS output:
$ astconvertt M31_r.fits M31_g.fits blank -oeps
```

The file type of the output will be specified with the (possibly complete) file name given to the `--output` option, which can either be given on the command-line or in any of the

configuration files (see Section 4.2 [Configuration files], page 59). Note that if the output suffix is not recognized, it will default to plain text format, see Section 5.2.1 [Recognized file formats], page 87.

The order of multiple input files is important. After reading the input file(s) the number of color channels in all the inputs will be used to define which color space is being used for the outputs and how each color channel is interpreted. Note that one file might have more than one color channel (for example in the JPEG format). If there is one color channel the output is grayscale, if three input color channels are given they are respectively considered to be the red, green and blue color channels and if there are four color channels they are respectively considered to be cyan, magenta, yellow and black.

The value to `--output` (or `-o`) can be either a full file name or just the suffix of the desired output format. In the former case, that same name will be used for the output. In the latter case, the name of the output file will be set based on the automatic output guidelines, see Section 4.8 [Automatic output], page 77. Note that the suffix name can optionally start a `.` (dot), so for example `--output=.jpg` and `--output=jpg` are equivalent. See Section 5.2.1 [Recognized file formats], page 87,

Besides the common set of options explained in Section 4.1.2 [Common options], page 52, the options to ConvertType can be classified into input, output and flux related options. The majority of the options are to do with the flux range. Astronomical data usually have a very large dynamic range (difference between maximum and minimum value) and different subjects might be better demonstrated with a limited flux range.

Input:

`-h STR/INT`
`--hdu=STR/INT`

In ConvertType, it is possible to call the HDU option multiple times for the different input FITS files (corresponding to different color channels) in the same order that they are called on the command-line. Except for the fact that multiple calls are possible, this option is identical to the common `--hdu` in Section 4.1.2.1 [Input/Output options], page 52. The number of calls to this option cannot be less than the number of input FITS files, but if there are more, the extra HDUs will be ignored, note that they will be read in the order described in Section 4.2.2 [Configuration file precedence], page 60.

Output:

`-w FLT`
`--widthincm=FLT`

The width of the output in centimeters. This is only relevant for those formats that accept such a width (not plain text for example). For most digital purposes, the number of pixels is far more important than the value to this parameter because you can adjust the absolute width (in inches or centimeters) in your document preparation program.

`-b INT`
`--borderwidth=INT`

The width of the border to be put around the EPS and PDF outputs in units of PostScript points. There are 72 or 28.35 PostScript points in an inch or

centimeter respectively. In other words, there are roughly 3 PostScript points in every millimeter. If you are planning on adding a border, its significance is highly correlated with the value you give to the `--widthincm` parameter.

Unfortunately in the document structuring convention of the PostScript language, the "bounding box" has to be in units of PostScript points with no fractions allowed. So the border values only have to be specified in integers. To have a final border that is thinner than one PostScript point in your document, you can ask for a larger width in ConvertType and then scale down the output EPS or PDF file in your document preparation program. For example by setting `width` in your `includegraphics` command in TEX or LATEX. Since it is vector graphics, the changes of size have no effect on the quality of your output quality (pixels don't get different values).

`-x`
`--hex` Use Hexadecimal encoding in creating EPS output. By default the ASCII85 encoding is used which provides a much better compression ratio. When converted to PDF (or included in TEX or LATEX which is finally saved as a PDF file), an efficient binary encoding is used which is far more efficient than both of them. The choice of EPS encoding will thus have no effect on the final PDF.

So if you want to transfer your EPS files (for example if you want to submit your paper to arXiv or journals in PostScript), their storage might become important if you have large images or lots of small ones. By default ASCII85 encoding is used which offers a much better compression ratio (nearly 40 percent) compared to Hexadecimal encoding.

`-u INT`
`--quality=INT`
The quality (compression) of the output JPEG file with values from 0 to 100 (inclusive). For other formats the value to this option is ignored. Note that only in grayscale (when one input color channel is given) will this actually be the exact quality (each pixel will correspond to one input value). If it is in color mode, some degradation will occur. While the JPEG standard does support loss-less graphics, it is not commonly supported.

Flux range:

`-c STR`
`--chang=STR`
(`=STR`) Change pixel values with the following format `"from1:to1, from2:to2,..."`. This option is very useful in displaying labeled pixels (not actual data images which have noise) like segmentation maps. In labeled images, usually a group of pixels have a fixed integer value. With this option, you can manipulate the labels before the image is displayed to get a better output for print or to emphasize on a particular set of labels and ignore the rest. The labels in the images will be changed in the same order given. By default first the pixel values will be converted then the pixel values will be truncated (see `--fluxlow` and `--fluxhigh`).

You can use any number for the values irrespective of your final output, your given values are stored and used in the double precision floating point format.

So for example if your input image has labels from 1 to 20000 and you only want to display those with labels 957 and 11342 then you can run ConvertType with these options:

```
$ astconvertt --change=957:50000,11342:50001 --fluxlow=5e4 \
   --fluxhigh=1e5 segmentationmap.fits --output=jpg
```

While the output JPEG format is only 8 bit, this operation is done in an intermediate step which is stored in double precision floating point. The pixel values are converted to 8-bit after all operations on the input fluxes have been complete. By placing the value in double quotes you can use as many spaces as you like for better readability.

`-C`
`--changeaftertrunc`
Change pixel values (with `--change`) after truncation of the flux values, by default it is the opposite.

`-L FLT`
`--fluxlow=FLT`
The minimum flux (pixel value) to display in the output image, any pixel value below this value will be set to this value in the output. If the value to this option is the same as `--fluxmax`, then no flux truncation will be applied. Note that when multiple channels are given, this value is used for all the color channels.

`-H FLT`
`--fluxhigh=FLT`
The maximum flux (pixel value) to display in the output image, see `--fluxlow`.

`-m INT`
`--maxbyte=INT`
This is only used for the JPEG and EPS output formats which have an 8-bit space for each channel of each pixel. The maximum value in each pixel can therefore be $2^8 - 1 = 255$. With this option you can change (decrease) the maximum value. By doing so you will decrease the dynamic range. It can be useful if you plan to use those values for other purposes.

`-A INT`
`--flminbyte=INT`
If the lowest pixel value in the input channels is larger than the value to `--fluxlow`, then that input value will be redundant. In some situations it might be necessary to set the minimum byte value (0) to correspond to that flux even if the data do not reach that value. With this option you can do this. Note that if the minimum pixel value is smaller than `--fluxlow`, then this option is redundant.

`-B INT`
`--fhmaxbyte=INT`
See `--flminbyte`.

`-i`
`--invert` For 8-bit output types (JPEG, EPS, and PDF for example) the final value that is stored is inverted so white becomes black and vice versa. The reason for this

is that astronomical images usually have a very large area of blank sky in them. The result will be that a large are of the image will be black. Note that this behaviour is ideal for grayscale images, if you want a color image, the colors are going to be mixed up.

5.3 Table

Tables are the products of processing astronomical images and spectra. For example in Gnuastro, MakeCatalog will process the defined pixels over an object and produce a catalog (see Section 7.3 [MakeCatalog], page 175). For each identified object, MakeCatalog can print its position on the image or sky, its total brightness and many other information that is deducible from the given image. Each one of these properties is a column in its output catalog (or table) and for each input object, we have a row.

When there are only a small number of objects (rows) and not too many properties (columns), then a simple plain text file is mainly enough to store, transfer, or even use the produced data. However, to be more efficient in all these aspects, astronomers have defined the FITS binary table standard to store data in a binary (0 and 1) format, not plain text. This can offer major advantages in all those aspects: the file size will be greatly reduced and the reading and writing will be faster (because the RAM and CPU also work in binary).

The FITS standard also defines a standard for ASCII tables, where the data are stored in the human readable ASCII format, but within the FITS file structure. These are mainly useful for keeping ASCII data along with images and possibly binary data as multiple (conceptually related) extensions within a FITS file. The acceptable table formats are fully described in Section 4.5 [Tables], page 66.

Binary tables are not easily readable by human eyes. There is no fixed/unified standard on how the zero and ones should be interpretted. The Unix-like operating systems have flourished because of a simple fact: communication between the various tools is based on human readable characters[3]. So while the FITS table standards are very beneficial for the tools that recognize them, they are hard to use in the vast majority of available software. This creates limitations for their generic use.

'Table' is Gnuastro's solution to this problem. With Table, FITS tables (ASCII or binary) are directly accessible to the Unix-like operating systems power-users (those working the command-line or shell, see Section 1.6.1 [Command-line interface], page 7). With Table, a FITS table (in binary or ASCII formats) is only one command away from AWK (or any other tool you want to use). Just like a plain text file that you read with the `cat` command. You can pipe the output of Table into any other tool for higher-level processing, see the examples in Section 5.3.1 [Invoking Table], page 94, for some simple examples.

5.3.1 Invoking Table

Table will read/write, select, convert, or show the information of the columns in FITS ASCII table, FITS binary table and plain text table files, see Section 4.5 [Tables], page 66. Output

[3] In "The art of Unix programming", Eric Raymond makes this suggestion to programmers: "When you feel the urge to design a complex binary file format, or a complex binary application protocol, it is generally wise to lie down until the feeling passes.". This is a great book and strongly recommended, give it a look if you want to truly enjoy your work/life in this environment.

columns can also be determined by number or regular expression matching of column names, units, or comments. The executable name is `asttable` with the following general template

```
$ asttable [OPTION...] InputFile
```

One line examples:

```
## Get the table column information (name, data type, or units):
$ asttable bintab.fits --information

## Only print those columns which have a name starting with "MAG_":
$ asttable bintab.fits --column=/^MAG_/

## Only print the 2nd column, and the third column multiplied by 5:
$ asttable bintab.fits | awk '{print $2, 5*$3}'

## Only print rows with a value in the 10th column above 100000:
$ asttable bintab.fits | awk '$10>10e5 {print}'

## Sort the output columns by the third column, save output:
$ asttable bintab.fits | 'sort -k3 > output.txt

## Convert a plain text table to a binary FITS table:
$ asttable plaintext.txt --output=table.fits --tabletype=fits-binary
```

In the absence of selected columns, all the input file's columns will be output. If the specified output is a FITS file, the type of FITS table (binary or ASCII) will be determined from the `--tabletype` option. If the output is not a FITS file, it will be printed as a plain text table (with space characters between the columns). When the columns are accompanied by meta-data (like column name, units, or comments), this information will also printed in the plain text file before the table, as described in Section 4.5.2 [Gnuastro text table format], page 69.

For the full list of options common to all Gnuastro programs please see Section 4.1.2 [Common options], page 52. Options can also be stored in directory, user or system-wide configuration files to avoid repeating on the command-line, see Section 4.2 [Configuration files], page 59. Table does not follow Automatic output that is common in most Gnuastro programs, see Section 4.8 [Automatic output], page 77. Thus, in the absence of an output file, the selected columns will be printed on the command-line with no column information, ready for redirecting to other tools like AWK or sort, similar to the examples above.

`-i`
`--information`

Only print the column information in the specified table on the command-line and exit. Each column's information (number, name, units, data type, and comments) will be printed as a row on the command-line. Note that the FITS standard only requires the data type (see Section 4.4 [Numeric data types], page 64), and in plain text tables, no meta-data/information is mandatory. Gnuastro has its own convention in the comments of a plain text table to store and transfer this information as described in Section 4.5.2 [Gnuastro text table format], page 69.

This option will take precedence over the `--column` option, so when it is called along with requested columns, the latter will be ignored. This can be useful if you forget the identifier of a column after you have already typed some on the command-line. You can simply add a `-i` and run Table to see the whole list and remember. Then you can use the shell history (with the up arrow key on the keyboard), and retrieve the last command with all the previously typed columns present, delete `-i` and add the identifier you had forgot.

`-c STR/INT`
`--column=STR/INT`

Specify the columns to output, see Section 4.5.3 [Selecting table columns], page 71, for a thorough explanation on how the value to this option is interpretted. To select several output columns, this option can also be called any number times in one call to Table. The order of the output columns will be the same call order on the command-line.

This option is not mandatory, if no specific columns are requested, all the input table columns are output. When this option is called multiple times, it is possible to output one column more than once.

6 Data manipulation

Images are one of the major formats of data that is used in astronomy. The functions in this chapter explain the GNU Astronomy Utilities which are provided for their manipulation. For example cropping out a part of a larger image or convolving the image with a given kernel or applying a transformation to it.

6.1 Crop

Astronomical images are often very large, filled with thousands of galaxies. It often happens that you only want a section of the image, or you have a catalog of sources and you want to visually analyze them in small postage stamps. Crop is made to do all these things. When more than one crop is required, Crop will divide the crops between multiple threads to significantly reduce the run time.

Astronomical surveys are usually extremely large. So large in fact, that the whole survey will not fit into a reasonably sized file. Because of this, surveys usually cut the final image into separate tiles and store each tile in a file. For example the COSMOS survey's Hubble space telescope, ACS F814W image consists of 81 separate FITS images, with each one having a volume of 1.7 Giga bytes.

Even though the tile sizes are chosen to be large enough that too many galaxies/targets don't fall on the edges of the tiles, inevitably some do. So when you simply crop the image of such targets from one tile, you will miss a large area of the surrounding sky (which is essential in estimating the noise). Therefore in its WCS mode, Crop will stitch parts of the tiles that are relevant for a target (with the given width) from all the input images that cover that region into the output. Of course, the tiles have to be present in the list of input files.

Besides cropping postage stamps around certain coordinates, Crop can also crop arbitrary polygons from an image (or a set of tiles by stitching the relevant parts of different tiles within the polygon), see `--polygon` in Section 6.1.4 [Invoking Crop], page 101. Alternatively, it can crop out rectangular regions through the `--section` option from one image, see Section 6.1.2 [Crop section syntax], page 100.

6.1.1 Crop modes

In order to be comprehensive, intuitive, and easy to use, Crop has two ways to define crop region: 1) From its center and (square) side length, for example to generate postage stamps of a given catalog. 2) The vertices of the crop region, this can be useful for larger crops over many targets, for example to crop out a uniformly deep, or contiguous, region of a large survey. Irrespective of how the crop region is defined, both Image/pixel, and World coordinate system (WCS) coordinates are acceptable and all coordinates are read as floating point numbers (not integers, except for the `--section` option, see below). The coordinate standards used are the main modes of Crop. Here, the different ways to specify the crop region are discussed within each standard. For the full list options, please see Section 6.1.4 [Invoking Crop], page 101.

When the crop is defined by its center, the respective (integer) central pixel position will be found internally according to the FITS standard. To have this pixel positioned in the center of the cropped region, the final cropped region must have (in Image mode), or

will have (in WCS mode) an add number of pixels. Furthermore, when the crop is defined as by its center, Crop allows you to only keep crops what don't have any blank pixels in the vicinity of their center (your primary target). This can be very convenient when your input catalog/coordinates originated from another survey/filter which is not fully covered by your input image, to learn more about this feature, please see the description of the `--checkcenter` option in Section 6.1.4 [Invoking Crop], page 101.

Image coordinates

In image mode, Crop uses the pixel coordinates/positions (instead of world coordinates). In image mode, only one image may be input. The crop(s) can be defined in multiple ways as listed below.

Center of multiple crops (in a catalog)

The center of (possibly multiple) crops are read from a text file. A catalog can also contain WCS coordinates, so `--mode=img` is necessary to avoid ambiguity when all columns are specified. The `--xcol` and `--ycol` columns in the catalog are interpreted as the center of a square crop with a width of `--iwidth` pixels (an odd number). The columns can contain any floating point value. The value to `--output` option is seen as a directory which will host (the possibly multiple) separate crop files, see Section 6.1.4.2 [Crop output], page 107, for more. For a tutorial using this feature, please see Section 2.1 [Hubble visually checks and classifies his catalog], page 14.

Center of a single crop (on the command-line)

The center of the crop is given on the command-line with the `--xc` and `--yc` options. The crop width is specified by the `--iwidth` option. The given coordinates can be any floating point and the width has to be an odd integer, see the explanations above.

Vertices of a single crop

In Image mode there are two options to define the vertices of a region to crop: `--section` and `--polygon`. The former is lower-level (doesn't accept floating point vertices, and only a rectangular region can be defined), it is also only available in Image mode. Please see Section 6.1.2 [Crop section syntax], page 100, for a full description of this method.

The latter option (`--polygon`) is a higher-level method to define any convex polygon (with any number of vertices) with floating point values. Please see the description of this option in Section 6.1.4 [Invoking Crop], page 101, for its syntax. It is available in both Image and WCS modes, hence, to read the vertices as pixel coordinates, `--mode=img` has to be called.

WCS coordinates

Use the World Coordinate System (WCS) information to define the crop, not pixel coordinates. In this mode, the width (`--wwidth`) is read in units of arcseconds and multiple images (tiles in a survey) can be input. When the cropped

region (defined by center or vertices) overlaps with multiple of the input images/tiles, the overlapping regions will be taken from the respective input (they will be stitched in the crop).

In this mode, the input images do not necessarily have to be the same size, they just need to have the same orientation and pixel resolution. Currently only orientation along the celestial coordinates is accepted, if your input has a different orientation you can use Warp's `--align` option to align the image before cropping it (see Section 6.4 [Warp], page 138).

Each individual input image/tile can even be smaller than the final crop. In any case, any part of any of the input images which overlaps with the desired region will be used in the crop. Note that if there is an overlap in the input images/tiles, the pixels from the last input image read are going to be used for the overlap. Crop will not change pixel values, so it assumes your overlapping tiles were cutout from the same original image. There are multiple ways to define your cropped region as listed below.

Center of multiple crops (in a catalog)
: Similar to catalog inputs in Image mode (above), but `--mode=wcs` has to be activated to avoid ambiguity. The central RA and Dec value for each crop will be read from the `--racol` and `--deccol` columns of the input catalog. The square cropped box will have an odd number of pixels.

Center of a single crop (on the command-line)
: You can specify the center of only one crop box with the `--ra` and `--dec` options. If it exists in the input images, it will be cropped similar to the catalog mode.

Vertices of a single crop
: The `--polygon` option is a high-level method to define any convex polygon (with any number of vertices). Please see the description of this option in Section 6.1.4 [Invoking Crop], page 101, for its syntax. It is available in both Image and WCS modes, hence, to read the vertices as RA and Dec coordinates, `--mode=wcs` has to be called/activated.

> **CAUTION:** In WCS mode, the image has to be aligned with the celestial coordinates, such that the first FITS axis is parallel (opposite direction) to the Right Ascension (RA) and the second FITS axis is parallel to the declination. If these conditions aren't met for an image, Crop will warn you and abort. You can use Warp's `--align` option to align the input image with these coordinates, see Section 6.4 [Warp], page 138.

As a summary, if you don't specify a catalog, you have to define the cropped region manually on the command-line. Other than the `--polygon` option, the other methods to define a single crop are uniqe to each mode, so the mode (`--mode`) will be ignored. When using a catalog to define the crop centers, if the columns to only one mode are given (for

example only `--xcol` and `--ycol`, not the WCS columns) then the mode is irrelevant. However, when all image and WCS mode columns are given, `--mode` is mandatory (you can keep the columns in a configuration file and simply set the mode, see Section 4.2 [Configuration files], page 59, and Section 4.5.3 [Selecting table columns], page 71). When using `--polygon` the mode is mandatory to interpret the values correctly.

6.1.2 Crop section syntax

When in image mode, one of the methods to crop only one rectangular section from the input image is to use the `--section` option. Crop has a powerful syntax to read the box parameters from a string of characters. If you leave certain parts of the string to be empty, Crop can fill them for you based on the input image sizes.

To define a box, you need the coordinates of two points: the first (`X1`, `Y1`) and the last pixel (`X2`, `Y2`) pixel positions in the image, or four integer numbers in total. The four coordinates can be specified with one string in this format: '`X1:X2,Y1:Y2`'. This string is given to the `--section` option. Therefore, the pixels along the first axis that are ≥`X1` and ≤`X2` will be included in the cropped image. The same goes for the second axis. Note that each different term will be read as an integer, not a float. This is a low-level option, for a higher-level way to specify region (any polygon, not just a box), please see the `--polygon` option in Section 6.1.4.1 [Crop options], page 102. Also note that in the FITS standard, pixel indexes along each axis start from unity(1) not zero(0).

You can omit any of the values and they will be filled automatically. The left hand side of the colon (`:`) will be filled with `1`, and the right side with the image size. So, `2:,:` will include the full range of pixels along the second axis and only those with a first axis index larger than `2` in the first axis. If the colon is omitted for a dimension, then the full range is automatically used. So the same string is also equal to `2:,` or `2:` or even `2`. If you want such a case for the second axis, you should set it to: `,2`.

If you specify a negative value, it will be seen as before the indexes of the image which are outside the image along the bottom or left sides when viewed in SAO ds9. In case you want to count from the top or right sides of the image, you can use an asterisk (`*`). When confronted with a `*`, Crop will replace it with the maximum length of the image in that dimension. So `*-10:*+10,*-20:*+20` will mean that the crop box will be 20×40 pixels in size and only include the top corner of the input image with 3/4 of the image being covered by blank pixels, see Section 6.1.3 [Blank pixels], page 100.

If you feel more comfortable with space characters between the values, you can use as many space characters as you wish, just be careful to put your value in double quotes, for example `--section="5:200, 123:854"`. If you forget the quotes, anything after the first space will not be seen by `--section` and you will most probably get an error because the rest of your string will be read as a filename (which most probably doesn't exist). See Section 4.1 [Command-line], page 48, for a description of how the command-line works.

6.1.3 Blank pixels

The cropped box can potentially include pixels that are beyond the image range. For example when a target in the input catalog was very near the edge of the input image. The parts of the cropped image that were not in the input image will be filled with the following two values depending on the data type of the image. In both cases, SAO ds9 will not color code those pixels.

- If the data type of the image is a floating point type (float or double), IEEE NaN (Not a number) will be used.
- For integer types, pixels out of the image will be filled with the value of the `BLANK` keyword in the cropped image header. The value assigned to it is the lowest value possible for that type, so you will probably never need it any way. Only for the unsigned character type (`BITPIX=8` in the FITS header), the maximum value is used because it is unsigned, the smallest value is zero which is often meaningful.

You can ask for such blank regions to not be included in the output crop image using the `--noblank` option. In such cases, there is no guarantee that the image size of your outputs are what you asked for.

In some survey images, unfortunately they do not use the `BLANK` FITS keyword. Instead they just give all pixels outside of the survey area a value of zero. So by default, when dealing with float or double image types, any values that are 0.0 are also regarded as blank regions. This can be turned off with the `--zeroisnotblank` option.

6.1.4 Invoking Crop

Crop will crop a region from an image. If in WCS mode, it will also stitch parts from separate images in the input files. The executable name is `astcrop` with the following general template

```
$ astcrop [OPTION...] [ASCIIcatalog] ASTRdata ...
```

One line examples:

```
## Crop all objects in catalog.txt from image.fits:
$ astcrop catalog.txt image.fits

## Crop all options in catalog (with RA,DEC) from all the files
## ending in '_drz.fits' in '/mnt/data/COSMOS/':
$ astcrop --mode=wcs catalog.txt /mnt/data/COSMOS/*_drz.fits

## Crop-out the outer 10 border pixels of input image:
$ astcrop --section=10:*-10,10:*-10 --hdu=2 image.fits

## Crop region around RA and Dec of (189.16704, 62.218203):
$ astcrop --ra=189.16704 --dec=62.218203 goodsnorth.fits

## Crop region around pixel coordinate (568.342, 2091.719):
$ astcrop --xc=568.342 --yc=2091.719 --iwidth=201 image.fits
```

Crop has one mandatory argument which is the input image name(s), shown above with `ASTRdata ...`. You can use shell expansions, for example `*` for this if you have lots of images in WCS mode. If the crop box centers are in a catalog, you also have to provide the catalog name as an argument. Alternatively, you have to provide the crop box parameters with command-line options. See Section 6.1.4.2 [Crop output], page 107, for how the output file name(s) can be specified. For the full list of general options to all Gnuastro programs (including Crop), please see Section 4.1.2 [Common options], page 52.

Floating point numbers can be used to specify the crop region (except the `--section` option, see Section 6.1.2 [Crop section syntax], page 100). In such cases, the floating point

values will be used to find the desired integer pixel indices based on the FITS standard. Hence, Crop ultimately doesn't do any sub-pixel cropping (in other words, it doesn't change pixel values). If you need such crops, you can use Section 6.4 [Warp], page 138, to first warp the image to the a new pixel grid where your initial floating points can be seen as integers, then crop from that with Crop. For example, let's say you want a crop from pixels 12.982 to 80.982 along the first dimension. You should first translate the image by −0.482 (note that the edge of a pixel is at integer multiples of 0.5). So you should run Warp with `--translate=-0.482,0` and then crop the warped image with `--section=13:81`.

There are two ways to define the cropped region: with its center or its vertices. See Section 6.1.1 [Crop modes], page 97, for a full description. In the former case, Crop can check if the central region of the cropped image is indeed filled with data or is blank (see Section 6.1.3 [Blank pixels], page 100), and not produce any output when the center is blank, see the description under `--checkcenter` for more.

When in catalog mode, Crop will run in parallel unless you set `--numthreads=1`, see Section 4.3 [Multi-threaded operations], page 62. Note that when multiple threads are being used, in verbose mode, the outputs will not be in order. This is because the threads are asynchronous and thus not started in order. This has no effect on the output, see Section 2.1 [Hubble visually checks and classifies his catalog], page 14, for a tutorial on effectively using this feature.

6.1.4.1 Crop options

The options can be classified into the following contexts: Input, Output and operating mode options. Options that are common to all Gnuastro program are listed in Section 4.1.2 [Common options], page 52, and will not be repeated here.

When you are specifying the crop vertices your self (through `--section`, or `--polygon`) on relatively small regions (depending on the resolution of your images) the outputs from image and WCS mode can be approximately equivalent. However, as the crop sizes get large, the curved nature of the WCS coordinates have to be considered. For example, when using `--section`, the right ascension of the bottom left and top left corners will not be equal. If you only want regions within a given right ascension, use `--polygon` in WCS mode.

> **NOTE:** The coordinates are in the FITS format. So the first axis is the horizontal axis when viewed in SAO ds9 and the second axis is the vertical. Also in the FITS standard, counting begins from 1 (one) not 0 (zero).

Input image parameters:

`--hstartwcs=INT`
: Specify the first keyword card (line number) to start finding the input image world coordinate system information. Distortions were only recently included in WCSLIB (from version 5). Therefore until now, different telescope would apply their own specific set of WCS keywords and put them into the image header along with those that WCSLIB does recognize. So now that WCSLIB recognizes most of the standard distortion parameters, they will get confused with the old

ones and give completely wrong results. For example in the CANDELS-GOODS South images[1].

The two `--hstartwcs` and `--hendwcs` are thus provided so when using older datasets, you can specify what region in the FITS headers you want to use to read the WCS keywords. Note that this is only relevant for reading the WCS information, basic data information like the image size are read separately. These two options will only be considered when the value to `--hendwcs` is larger than that of `--hstartwcs`. So if they are equal or `--hstartwcs` is larger than `--hendwcs`, then all the input keywords will be parsed to get the WCS information of the image.

`--hendwcs=INT`
: Specify the last keyword card to read for specifying the image world coordinate system on the input images. See `--hstartwcs`

Crop box parameters:

`-x FLT`
`--xc=FLT`
: X axis (first image axis, horizontal when viewed in SAO ds9) position of crop box center. Along with `--yc`, this only produces one crop in each run. The width of the square crop (in pixels) is the value to `--iwidth`.

`-y FLT`
`--yc=FLT`
: Y axis (second image axis, vertical when viewed in SAO ds9) position of crop box center. See `--xc`.

`-r FLT`
`--ra=FLT`
: Right Ascension of crop box center. Along with `--dec`, this only produces one crop in each run. The width of the square crop (in arcseconds) is the value to `--wwidth`.

`-d FLT`
`--dec=FLT`
: Declination of crop box center. see `--ra`.

`-a INT`
`--iwidth=INT`
: Width of the cropped region in units of pixels when the crop is defined by its center in Image mode (see Section 6.1.1 [Crop modes], page 97). In order for the chosen central pixel to be in the center of the cropped image, this option only accepts an odd number (an error will be given if its even). If you want an even sided crop, you can run Crop afterwards with `--section=":*-1,:*-1"` or `--section=2:,2:` (depending on which side you don't need), see Section 6.1.2 [Crop section syntax], page 100.

`-w FLT`
`--wwidth=FLT`
: The width of the crop box in WCS mode in units of arc-seconds.

---
[1] `https://archive.stsci.edu/pub/hlsp/candels/goods-s/gs-tot/v1.0/`

`-l STR`
`--polygon=STR`

String of crop polygon vertices. Note that currently only convex polygons should be used. In the future we will make it work for all kinds of polygons. Convex polygons are polygons that do not have an internal angle more than 180 degrees. This option can be used both in the image and WCS modes, see Section 6.1.1 [Crop modes], page 97. The cropped image will be the size of the rectangular region that completely encompasses the polygon. By default all the pixels that are outside of the polygon will be set as blank values (see Section 6.1.3 [Blank pixels], page 100). However, if `--outpolygon` is called all pixels interal to the vertices will be set to blank.

The syntax for the polygon vertices is similar to, and simpler than, that for `--section`. In short, the dimensions of each coordinate are separated by a comma (`,`) and each vertice is separated by a colon (`:`). You can define as many vertices as you like. If you would like to use space characters between the dimensions and vertices to make them more human-readable, then you have to put the value to this option in double quotation marks.

For example, let's assume you want to work on the deepest part of the WFC3/IR images of Hubble Space Telescope eXtreme Deep Field (HST-XDF). According to the webpage (`https://archive.stsci.edu/prepds/xdf/`)[2] the deepest part is contained within the coordinates:

```
[ (53.187414,-27.779152), (53.159507,-27.759633),
  (53.134517,-27.787144), (53.161906,-27.807208) ]
```

They have provided mask images with only these pixels in the WFC3/IR images, but what if you also need to work on the same region in the full resolution ACS images? Also what if you want to use the CANDELS data for the shallow region? Running Crop with `--polygon` will easily pull out this region of the image for you irrespective of the resolution. If you have set the operating mode to WCS mode in your nearest configuration file (see Section 4.2 [Configuration files], page 59), there is no need to call `--mode=wcs` on the command line. You may also provide many FITS images/tiles and Crop will stitch them to produce this cropped region:

```
$ astcrop --mode=wcs desired-filter-image(s).fits           \
   --polygon="53.187414,-27.779152 : 53.159507,-27.759633 : \
              53.134517,-27.787144 : 53.161906,-27.807208"
```

In other cases, you have an image and want to define the polygon yourself (it isn't already published like the example above). As the number of vertices increases, checking the vertice coordinates on a FITS viewer (for example SAO ds9) and typing them in one by one can be very tedious and prone to typo errors.

You can take the following steps to avoid the frustration and possible typos: Open the image with ds9 and activate its "region" mode with Edit→Region. Then define the region as a polygon with Region→Shape→Polygon. Click on the approximate center of the region you want and a small square will appear.

---

[2] `https://archive.stsci.edu/prepds/xdf/`

By clicking on the vertices of the square you can shrink or expand it, clicking and dragging anywhere on the edges will enable you to define a new vertice. After the region has been nicely defined, save it as a file with Region→Save Regions. You can then select the name and address of the output file, keep the format as `REG` and press "OK". In the next window, keep format as "ds9" and "Coordinate System" as "fk5". A plain text file (let's call it `ds9.reg`) is now created.

You can now convert this plain text file to Crop's polygon format with this command (when typing on the command-line, ignore the "\" at the end of the first and second lines along with the extra spaces, these are only for nice printing):

```
$ v=$(awk 'NR==4' ds9.reg | sed -e's/polygon(//'        \
           -e's/\([^,]*,[^,]*\),/\1:/g' -e's/)//' )
$ astcrop --mode=wcs image.fits --polygon=$v
```

`--outpolygon`

Keep all the regions outside the polygon and mask the inner ones with blank pixels (see Section 6.1.3 [Blank pixels], page 100). This is practically the inverse of the default mode of treating polygons. Note that this option only works when you have only provided one input image. If multiple images are given (in WCS mode), then the full area covered by all the images has to be shown and the polygon excluded. This can lead to a very large area if large surveys like COSMOS are used. So Crop will abort and notify you. In such cases, it is best to crop out the larger region you want, then mask the smaller region with this option.

`-s STR`
`--section=STR`

Section of the input image which you want to be cropped. See Section 6.1.2 [Crop section syntax], page 100, for a complete explanation on the syntax required for this input.

`-i STR/INT`
`--xcol=STR/INT`

Column selection for the first FITS axis position of the box center in a catalog. In SAO ds9, the first FITS axis is the horizontal axis. The value can be either the column number (starting from 1), or a match/search in the table meta-data, see Section 4.5.3 [Selecting table columns], page 71.

`-j STR/INT`
`--ycol=STR/INT`

Column selection for the second FITS axis position of the box center. In SAO ds9, the second FITS axis is the vertical axis. The value can be either the column number (starting from 1), or a match/search in the table meta-data, see Section 4.5.3 [Selecting table columns], page 71.

`-f STR/INT`
`--racol=STR/INT`
: Column selection for Right Ascension (RA) in the input catalog. The value can be either the column number (starting from 1), or a match/search in the table meta-data, see Section 4.5.3 [Selecting table columns], page 71.

`-g STR/INT`
`--deccol=STR/INT`
: Column selection of declination in the input catalog. The value can be either the column number (starting from 1), or a match/search in the table meta-data, see Section 4.5.3 [Selecting table columns], page 71.

`-n STR/INT`
`--namecol=STR/INT`
: Column selection of crop file name. The value can be either the column number (starting from 1), or a match/search in the table meta-data, see Section 4.5.3 [Selecting table columns], page 71. This option can be used both in Image and WCS modes, and not a mandatory. When a column is given to this option, the final crop base file name will be taken from the contents of this column. The directory will be determined by the `--output` option (current directory if not given) and the value to `--suffix` will be appended. When this column isn't given, the row number will be used instead.

Output options:

`-c INT`
`--checkcenter=INT`
: Box width (odd number of pixels) of region in the center of the image to check for blank values. If the value to this option is zero, no checking is done. This option is only relevant when the cropped region(s) are defined by their center (not by the vertices, see Section 6.1.1 [Crop modes], page 97). If any of the pixels in this central region of a crop (defined by its center) are blank, then it will not be created.

  Because survey regions don't often have a clean square or rectangle shape, some of the pixels on the sides of the survey FITS image don't commonly have any data and are blank (see Section 6.1.3 [Blank pixels], page 100). So when the catalog was not generated from the input image, it often happens that the image does not have data over some of the points.

  When the given center of a crop falls in such regions and this option has a non-zero, odd value, no crop will be created. Therefore with this option, you can specify a width of a small box (3 pixels is often good enough) around the central pixel of the cropped image. You can check which crops were created and which weren't from the command-line (if `--quiet` was not called, see Section 4.1.2.3 [Operating mode options], page 55), or in Crop's log file (see Section 6.1.4.2 [Crop output], page 107).

`-p STR`
`--suffix=STR`
: The suffix (or post-fix) of the output files for when you want all the cropped images to have a special ending. One case where this might be helpful is when

besides the science images, you want the weight images (or exposure maps, which are also distributed with survey images) of the cropped regions too. So in one run, you can set the input images to the science images and `--suffix=_s.fits`. In the next run you can set the weight images as input and `--suffix=_w.fits`.

`-b`
`--noblank`
: Pixels outside of the input image that are in the crop box will not be used. By default they are filled with blank values (depending on type), see Section 6.1.3 [Blank pixels], page 100. This option only applies only in Image mode, see Section 6.1.1 [Crop modes], page 97.

`-z`
`--zeroisnotblank`
: In float or double images, it is common to give the value of zero to blank pixels. If the input image type is one of these two types, such pixels will also be considered as blank. You can disable this behavior with this option, see Section 6.1.3 [Blank pixels], page 100.

Operating mode options:

`-O STR`
`--mode=STR`
: Operate in Image mode or WCS mode when the input coordinates can be both image or WCS. The value must either be `img` or `wcs`, see Section 6.1.1 [Crop modes], page 97, for a full description.

6.1.4.2 Crop output

The value given to `--output` option will depend on how many crops were created, see Section 6.1.1 [Crop modes], page 97:

- When a catalog is given, the value of the `--output` (see Section 4.1.2 [Common options], page 52) will be read as the directory to store the output cropped images. Hence if it doesn't alread exist, Crop will abort with an error of a "No such file or directory" error. The crop file names will consist of two parts: a variable part (the row number of each target starting from 1) along with a fixed string which you can set with the `--suffix` option.
- When only one crop is desired, the value to `--output` will be read as a file name. If no output is specified or if it is a directory, the output file name will follow the automatic output names of Gnuastro, see Section 4.8 [Automatic output], page 77: The string given to `--suffix` will be replaced with the `.fits` suffix of the input.

The header of each output cropped image will contain the names of the input image(s) it was cut from. If a name is longer than the 70 character space that the FITS standard allows for header keyword values, the name will be cut into several keywords from the nearest slash (`/`). The keywords have the following format: `ICFn_m` (for Crop File). Where `n` is the number of the image used in this crop and `m` is the part of the name (it can be broken into multiple keywords). Following the name is another keyword named `ICFnPIX` which shows the pixel range from that input image in the same syntax as Section 6.1.2

[Crop section syntax], page 100. So this string can be directly given to the `--section` option later.

Once done, a log file will be created in the current directory named `astcrop.log`. This file will have three columns and the same number of rows as the number of cropped images.

1. The cropped image file name for that row.
2. The number of input images that were used to create that image.
3. A `0` if the central few pixels (value to the `--checkcenter` option) are blank and `1` if they aren't. When the crop was not defined by its center (see Section 6.1.1 [Crop modes], page 97), or `--checkcenter` was given a value of 0 (see Section 6.1.4 [Invoking Crop], page 101), the center will not be checked and this column will be given a value of `-1`.

There are also comments on the top of the log file explaining basic information about the run and descriptions for the columns. If the log file cannot be created (for example you don't have write permission in the directory you are running Crop in) or you have specifically asked for no log file (with the `--nolog` option), then a log file will not be created (unless `--individual` is called).

6.2 Arithmetic

It is commonly necessary to do operations on some or all of the elements of a dataset independently (pixels in an image). For example, in the reduction of raw data it is necessary to subtract the Sky value (Section 7.1.3 [Sky value], page 150) from each image image. Later (once the images as warped into a single grid using Warp for example, see Section 6.4 [Warp], page 138), the images are co-added (the output pixel grid is the average of the pixels of the individual input images). Arithmetic is Gnuastro's program for such operations on your datasets directly from the command-line. It currently uses the reverse polish or postfix notation, see Section 6.2.1 [Reverse polish notation], page 108, and will work on the native data types of the input images/data to reduce CPU and RAM resources, see Section 4.4 [Numeric data types], page 64. For more information on how to run Arithmetic, please see Section 6.2.3 [Invoking Arithmetic], page 114.

6.2.1 Reverse polish notation

The most common notation for arithmetic operations is the infix notation (`https://en.wikipedia.org/wiki/Infix_notation`) where the operator goes between the two operands, for example $4 + 5$. While the infix notation is the preferred way in most programming languages, currently Arithmetic does not use it since it will require parenthesis which can complicate the implementation of the code. In the near future we do plan to adopt this notation[3], but for the time being (due to time constraints on the developers), Arithmetic uses the postfix or reverse polish notation (`https://en.wikipedia.org/wiki/Reverse_Polish_notation`). The Wikipedia article provides some excellent explanation on this notation but here we will give a short summary here for self-sufficiency.

In the postfix notation, the operator is placed after the operands, as we will see below this removes the need to define parenthesis for most ordinary operators. For example, instead of writing `5+6`, we write `5 6 +`. To easily understand how this notation works, you

[3] `https://savannah.gnu.org/task/index.php?13867`

can think of each operand as a node in a first-in-first-out stack. Every time an operator is confronted, it pops the number of operands it needs from the top of the stack (so they don't exist in the stack any more), does its operation and pushes the result back on top of the stack. So if you want the average of 5 and 6, you would write: `5 6 + 2 /`. The operations that are done are:

1. `5` is an operand, so it is pushed to the top of the stack.
2. `6` is an operand, so it is pushed to the top of the stack.
3. `+` is a binary operator, so pull the top two elements of the stack and perform addition on them (the order is $5 + 6$ in the example above). The result is `11`, push it on top of the stack.
4. `2` is an operand so push it onto the top of the stack.
5. `/` is a binary operator, so pull out the top two elements of the stack (top-most is `2`, then `11`) and divide the second one by the first.

In the Arithmetic program, the operands can be FITS images or numbers. As you can see, very complicated procedures can be created without the need for parenthesis or worrying about precedence. Even functions which take an arbitrary number of arguments can be defined in this notation. This is a very powerfull notation and is used in languages like Postscript[4] (the programming language in Postscript and compiled into PDF files) uses this notation.

6.2.2 Arithmetic operators

The recognized operators in Arithmetic are listed below. See Section 6.2.1 [Reverse polish notation], page 108, for more on how the operators and operands should be ordered on the command-line. The operands to all operators can be a data array (for example a FITS image) or a number, the output will be an array or number according to the inputs. For example a number multiplied by an array will produce an array. The conditional operators will return pixel, or numerical values of 0 (false) or 1 (true) and stored in an `unsigned char` data type (see Section 4.4 [Numeric data types], page 64).

`+` Addition, so "`4 5 +`" is equivalent to $4 + 5$.

`-` Subtraction, so "`4 5 -`" is equivalent to $4 - 5$.

`x` Multiplication, so "`4 5 x`" is equivalent to 4×5.

`/` Division, so "`4 5 /`" is equivalent to $4/5$.

`%` Modulo (remainder), so "`3 2 %`" is equivalent to 1. Note that the modulo operator only works on integer types.

`abs` Absolute value of first operand, so "`4 abs`" is equivalent to $|4|$.

`pow` First operand to the power of the second, so "`4.3 5f pow`" is equivalent to 4.3^5. Currently `pow` will only work on single or double precision floating point numbers or images. To be sure that a number is read as a floating point (even if it doesn't have any non-zero decimals) put an `f` after it.

---

[4] See the EPS and PDF part of Section 5.2.1 [Recognized file formats], page 87, for a little more on the Postscript language.

`sqrt` The square root of the first operand, so “`5 sqrt`” is equivalent to $\sqrt{5}$. The output type is determined from the input, so the output of this example will be `2` (since `5` doesn’t have any non-zero decimal digits). If you want `2.23607`, run `5f sqrt` instead, the `f` will ensure that a number will be read as a floating point number, even if it doesn’t have decimal digits. If the input image has an integer type, you should explicitly convert the image to floating point, for example `a.fits float sqrt`, see the type conversion operators below.

`log` Natural logarithm of first operand, so “`4 log`” is equivalent to $\ln(4)$. The output type is determined from the input, see the explanation under `sqrt` for more.

`log10` Base-10 logarithm of first operand, so “`4 log10`” is equivalent to $\log(4)$. The output type is determined from the input, see the explanation under `sqrt` for more.

`minvalue` Minimum (non-blank) value in the top operand on the stack, so “`a.fits minvalue`” will push the the minimum pixel value in this image onto the stack. Therefore this operator is mainly intended for data (for example images), if the top operand is a number, this operator just returns it without any change. So note that when this operator acts on a single image, the output will no longer be an image, but a number. The output of this operand is in the same type as the input.

`maxvalue` Maximum (non-blank) value of first operand in the same type, similar to `minvalue`.

`numvalue` Number of non-blank elements in first operand in the `uint64` type, similar to `minvalue`.

`sumvalue` Sum of non-blank elements in first operand in the `float32` type, similar to `minvalue`.

`meanvalue`
Mean value of non-blank elements in first operand in the `float32` type, similar to `minvalue`.

`stdvalue` Standard deviation of non-blank elements in first operand in the `float32` type, similar to `minvalue`.

`medianvalue`
Median of non-blank elements in first operand with the same type, similar to `minvalue`.

`min` The first popped operand to this operator must be a positive integer number which specifies how many further operands should be popped from the stack. The given number of operands must have the same type and size. Each pixel of the output of this operator will be set to the minimum value of the given number of operands (images) in that pixel.

For example the following command will produce an image with the same size and type as the inputs but each output pixel is set to the minimum respective pixel value of the three input images.

```
$ astarithmetic a.fits b.fits c.fits 3 min
```

Important notes:

- NaN/blank pixels will be ignored, see Section 6.1.3 [Blank pixels], page 100.
- The output will have the same type as the inputs. This is natural for the `min` and `max` operators, but for other similar operators (for example `sum`, or `average`) the per-pixel operations will be done in double precision floating point and then stored back in the input type. Therefore, if the input was an integer, C's internal type conversion will be used.

`max` Similar to `min`, but the pixels of the output will contain the maximum of the respective pixels in all operands in the stack.

`num` Similar to `min`, but the pixels of the output will contain the number of the respective non-blank pixels in all input operands.

`sum` Similar to `min`, but the pixels of the output will contain the sum of the respective pixels in all input operands.

`mean` Similar to `min`, but the pixels of the output will contain the mean (average) of the respective pixels in all operands in the stack.

`std` Similar to `min`, but the pixels of the output will contain the standard deviation of the respective pixels in all operands in the stack.

`median` Similar to `min`, but the pixels of the output will contain the median of the respective pixels in all operands in the stack.

`lt` Less than: If the second popped (or left operand in infix notation, see Section 6.2.1 [Reverse polish notation], page 108) value is smaller than the first popped operand, then this function will return a value of 1, otherwise it will return a value of 0. If both operands are images, then all the pixels will be compared with their counterparts in the other image. If only one operand is an image, then all the pixels will be compared with the the single value (number) of the other operand. Finally if both are numbers, then the output is also just one number (0 or 1). When the output is not a single number, it will be stored as an `unsigned char` type.

`le` Less or equal: similar to `lt` ('less than' operator), but returning 1 when the second popped operand is smaller or equal to the first.

`gt` Greater than: similar to `lt` ('less than' operator), but returning 1 when the second popped operand is greater than the first.

`ge` Greater or equal: similar to `lt` ('less than' operator), but returning 1 when the second popped operand is larger or equal to the first.

`eq` Equality: similar to `lt` ('less than' operator), but returning 1 when the two popped operands are equal (to double precision floating point accuracy).

`ne` Non-Equality: similar to `lt` ('less than' operator), but returning 1 when the two popped operands are *not* equal (to double precision floating point accuracy).

`and` Logical AND: returns 1 if both operands have a non-zero value and 0 if both are zero. Both operands have to be the same kind: either both images or both numbers.

`or` Logical OR: returns 1 if either one of the operands is non-zero and 0 only when both operators are zero. Both operands have to be the same kind: either both images or both numbers.

`not` Logical NOT: returns 1 when the operand is zero and 0 when the operand is non-zero. The operand can be an image or number, for an image, it is applied to each pixel separately.

`isblank` Test for a blank value (see Section 6.1.3 [Blank pixels], page 100). In essence, this is very similar to the conditional operators: the output is either 1 or 0 (see the 'less than' operator above). The difference is that it only needs one operand. Because of the definition of a blank pixel, a blank value is not even equal to itself, so you cannot use the equal operator above to select blank pixels. See the "Blank pixels" box below for more on Blank pixels in Arithmetic.

`where` Change the input (pixel) value *where*/if a certain condition holds. The conditional operators above can be used to define the condition. Three operands are required for `where`. The input format is demonstrated in this simplified example:

```
$ astarithmetic modify.fits binary.fits if-true.fits where
```

The value of any pixel in `modify.fits` that corresponds to a non-zero pixel of `binary.fits` will be changed to the value of the same pixel in `if-true.fits` (this may also be a number). The 3rd and 2nd popped operands (`modify.fits` and `binary.fits` respectively, see Section 6.2.1 [Reverse polish notation], page 108) have to have the same dimensions/size. `if-true.fits` can be either a number, or have the same dimension/size as the other two.

The 2nd popped operand (`binary.fits`) has to have `uint8` (or `unsigned char` in standard C) type (see Section 4.4 [Numeric data types], page 64). It is treated as a binary dataset (with only two values: zero and non-zero, hence the name `binary.fits` in this example). However, commonly you won't be dealing with an actual FITS file of a condition/binary image. You will probably define the condition in the same run based on some other reference image and use the conditional and logical operators above to make a true/false (or one/zero) image for you internally. For example the case below:

```
$ astarithmetic in.fits reference.fits 100 gt new.fits where
```

In the example above, any of the `in.fits` pixels that has a value in `reference.fits` greater than `100`, will be replaced with the corresponding pixel in `new.fits`. Effectively the `reference.fits 100 gt` part created the condition/binary image which was added to the stack (in memory) and later used by `where`. The command above is thus equivalent to these two commands:

```
$ astarithmetic reference.fits 100 gt --output=binary.fits
$ astarithmetic in.fits binary.fits new.fits where
```

Finally, the input operands are read and used independently, so you can use the same file more than once as any of the operands.

When the 1st popped operand to `where` (`if-true.fits`) is a single number, it may be a NaN value (or any blank value, depending on its type) like the example

below (see Section 6.1.3 [Blank pixels], page 100). When the number is blank, it will be converted to the blank value of the type of the 3rd popped operand (`in.fits`). Hence, in the example below, all the pixels in `reference.fits` that have a value greater than 100, will become blank in the natural data type of `in.fits` (even though NaN values are only defined for floating point types).

```
$ astarithmetic in.fits reference.fits 100 gt nan where
```

`bitand` Bitwise AND operator: only bits with values of 1 in both popped operands will get the value of 1, the rest will be set to 0. For example (assuming numbers can be written as bit strings on the command-line): `00101000 00100010 bitand` will give `00100000`. Note that the bitwise operators only work on integer type datasets.

`bitor` Bitwise inclusive OR operator: The bits where atleast one of the two popped operands has a 1 value get a value of 1, the others 0. For example (assuming numbers can be written as bit strings on the command-line): `00101000 00100010 bitand` will give `00101010`. Note that the bitwise operators only work on integer type datasets.

`bitxor` Bitwise exclusive OR operator: A bit will be 1 if it differs between the two popped operands. For example (assuming numbers can be written as bit strings on the command-line): `00101000 00100010 bitand` will give `00001010`. Note that the bitwise operators only work on integer type datasets.

`lshift` Bitwise left shift operator: shift all the bits of the first operand to the left by a number of times given by the second operand. For example (assuming numbers can be written as bit strings on the command-line): `00101000 2 lshift` will give `10100000`. This is equivalent to multiplication by 4. Note that the bitwise operators only work on integer type datasets.

`rshift` Bitwise right shift operator: shift all the bits of the first operand to the right by a number of times given by the second operand. For example (assuming numbers can be written as bit strings on the command-line): `00101000 2 rshift` will give `00001010`. Note that the bitwise operators only work on integer type datasets.

`bitnot` Bitwise not (more formally known as one's complement) operator: flip all the bits of the popped operand (note that this is the only unary, or single operand, bitwise operator). In other words, any bit with a value of `0` is changed to `1` and vice-versa. For example (assuming numbers can be written as bit strings on the command-line): `00101000 bitnot` will give `11010111`. Note that the bitwise operators only work on integer type datasets/numbers.

`uint8` Convert the type of the popped operand to 8-bit un-signed integer type (see Section 4.4 [Numeric data types], page 64). The internal conversion of C will be used.

`int8` Convert the type of the popped operand to 8-bit signed integer type (see Section 4.4 [Numeric data types], page 64). The internal conversion of C will be used.

`uint16` Convert the type of the popped operand to 16-bit un-signed integer type (see Section 4.4 [Numeric data types], page 64). The internal conversion of C will be used.

`int16` Convert the type of the popped operand to 16-bit signed integer (see Section 4.4 [Numeric data types], page 64). The internal conversion of C will be used.

`uint32` Convert the type of the popped operand to 32-bit un-signed integer type (see Section 4.4 [Numeric data types], page 64). The internal conversion of C will be used.

`int32` Convert the type of the popped operand to 32-bit signed integer type (see Section 4.4 [Numeric data types], page 64). The internal conversion of C will be used.

`uint64` Convert the type of the popped operand to 64-bit un-signed integer (see Section 4.4 [Numeric data types], page 64). The internal conversion of C will be used.

`float32` Convert the type of the popped operand to 32-bit (single precision) floating point (see Section 4.4 [Numeric data types], page 64). The internal conversion of C will be used.

`float64` Convert the type of the popped operand to 64-bit (double precision) floating point (see Section 4.4 [Numeric data types], page 64). The internal conversion of C will be used.

Blank pixels in Arithmetic: Blank pixels in the image (see Section 6.1.3 [Blank pixels], page 100) will be stored based on the data type. When the input is floating point type, blank values are NaN. One aspect of NaN values is that by definition they will fail on *any* comparison. Hence both equal and not-equal operators will fail when both their operands are NaN! Therefore, the only way to guarantee selection of blank pixels is through the `isblank` operator explained above.

One way you can exploit this property of the NaN value to your advantage is when you want a fully zero-valued image (even over the blank pixels) based on an already existing image (with same size and world coordinate system settings). The following command will produce this for you:

```
$ astarithmetic input.fits nan eq --output=all-zeros.fits
```

Note that on the command-line you can write NaN in any case (for example `NaN`, or `NAN` are also acceptable). Reading NaN as a floating point number in Gnuastro isn't case-sensitive.

6.2.3 Invoking Arithmetic

Arithmetic will do pixel to pixel arithmetic operations on the individual pixels of input data and/or numbers. For the full list of operators with explanations, please see Section 6.2.2 [Arithmetic operators], page 109. Any operand that only has a single element (number, or single pixel FITS image) will be read as a number, the rest of the inputs must have the same dimensions. The general template is:

```
$ astarithmetic [OPTION...] ASTRdata1 [ASTRdata2] OPERATOR ...
```

One line examples:

```
## Calculate (10.32-3.84)^2.7 quietly (will just print 155.329):
$ astarithmetic -q 10.32 3.84 - 2.7 pow

## Inverse the input image (1/pixel):
$ astarithmetic 1 image.fits / --out=inverse.fits

## Multiply each pixel in image by -1:
$ astarithmetic image.fits -1 x --out=negative.fits

## Subtract extension 4 from extension 1 (counting from zero):
$ astarithmetic image.fits image.fits - --out=skysub.fits           \
                --hdu=1 --hdu=4

## Add two images, then divide them by 2 (2 is read as floating point):
$ astarithmetic image1.fits image2.fits + 2f / --out=average.fits

## Use Arithmetic's average operator:
$ astarithmetic image1.fits image2.fits average --out=average.fits

## Calculate the median of three images in three separate extensions:
$ astarithmetic img1.fits img2.fits img3.fits median                \
                -h0 -h1 -h2 --out=median.fits
```

If the output is an image, and the `--output` option is not given, automatic output will use the name of the first FITS image encountered to generate an output file name, see Section 4.8 [Automatic output], page 77. Also, output WCS information will be taken from the first input image encountered. When the output is a single number, that number will be printed in the standard output and no output file will be created. Arithmetic's notation for giving operands to operators is described in Section 6.2.1 [Reverse polish notation], page 108. To ignore certain pixels, set them as blank, see Section 6.1.3 [Blank pixels], page 100, for example with the `where` operator (see Section 6.2.2 [Arithmetic operators], page 109). See Section 4.1.2 [Common options], page 52, for a review of the options in all Gnuastro programs. Arithmetic just redefines the `--hdu` option as explained below:

`-h INT/STR`
`--hdu INT/STR`

The header data unit of the input FITS images, see Section 4.1.2.1 [Input/Output options], page 52. Unlike most options in Gnuastro (which will ultimately only have one value for this option), Arithmetic allows `--hdu` to be called multiple times and the value of each invocation will be stored separately (for the unlimited number of input images you would like to use). Recall that for other programs this (common) option only takes a single value. So in other programs, if you specify it multiple times on the command-line, only the last value will be used and in the configuration files, it will be ignored if it already has a value.

The order of the values to `--hdu` has to be in the same order as input FITS images. Options are first read from the command-line (from left to right),

then top-down in each configuration file, see Section 4.2.2 [Configuration file precedence], page 60.

If the number of HDUs is less than the number of input images, Arithmetic will abort and notify you. However, if there are more HDUs than FITS images, there is no problem: they will be used in the given order (every time a FITS image comes up on the stack) and the extra HDUs will be ignored in the end. So there is no problem with having extra HDUs in the configuration files and by default several HDUs with a value of `0` are kept in the system-wide configuration file when you install Gnuastro.

Arithmetic accepts two kinds of input: images and numbers. Images are considered to be any of the inputs that is a file name of a recognized type (see Section 4.1.1.1 [Arguments], page 49) and has more than one element/pixel. Numbers on the command-line will be read into the smallest type (see Section 4.4 [Numeric data types], page 64) that can store them, so `-2` will be read as a `char` type (which is signed on most systems and can thus keep negative values), `2500` will be read as an `unsigned short` (all positive numbers will be read as unsigned), while `3.1415926535897` will be read as a `double` and `3.14` will be read as a `float`. To force a number to be read as float, add a `f` after it, so `5f` will be added to the stack as `float` (see Section 6.2.1 [Reverse polish notation], page 108).

Unless otherwise stated (in Section 6.2.2 [Arithmetic operators], page 109), the operators can deal with multiple datatypes. For example in "`a.fits b.fits +`", the image types can be `long` and `float`. In such cases, C's internal type conversion will be used. The output type will be set to the higher-ranking type of the two inputs. Unsigned integer types have smaller ranking than their signed counterparts and floating point types have higher ranking than the integer types. So the internal C type conversions done in the example above are equivalent to this piece of C:

```
size_t i;
long a[100];
float b[100], out;
for(i=0;i<100;++i) out[i]=a[i]+b[i];
```

Relying on the default C type conversion significantly speeds up the processing and also requires less RAM (when using very large images). However this great advantage comes at the cost of preparing for all the combinations of types while building/compiling Gnuastro. With the full list of CFITSIO types, compilation can take roughly half an hour. However, some types are not too common, therfore Gnuastro comes with a set of configure time options letting you enable certain types for native compilation. You can see the full list of `--enable-bin-op-XXXX` options in Section 3.3.1.1 [Gnuastro configure options], page 37.

When a type isn't enabled for native binary operations, the input data will be internally converted to the smallest, larger type that was enabled. This can slow down your processing (which is faster for smaller/integer types) and consume more RAM (to copy the new type), so if you often deal with data of a specific types, it is much better to make the one-time investment at compilation time and reap the benefits each time you run Gnuastro/Arithmetic. Note all arithmetic operations are done by `gal_data_arithmetic` function in Section 10.3 [Gnuastro library], page 233, so the choice of native binary operator types will affect any program (within Gnuastro or outside of it) that uses this function (including Arithmetic).

Some operators can only work on integer types (of any length, for example bitwise operators) while others only work on floating point types, (currently only the `pow` operator). In such cases, if the operand type(s) are different an error will be printed and internal conversion won't occur. Arithmetic also comes with internal type conversion operators which you can use to convert the data into the appropriate type, see Section 6.2.2 [Arithmetic operators], page 109.

The hyphen (`-`) can be used both to specify options (see Section 4.1.1.2 [Options], page 50) and also to specify a negative number which might be necessary in your arithmetic. In order to enable you to do this, Arithmetic will first parse all the input strings and if the first character after a hyphen is a digit, then that hyphen is temporarily replaced by the vertical tab character which is not commonly used. The arguments are then parsed and these strings will not be specified as an option. Then the given arguments are parsed and any vertical tabs are replaced back with a hyphen so they can be read as negative numbers. Therefore, as long as the names of the files you want to work on, don't start with a vertical tab followed by a digit, there is no problem. An important consequence of this implementation is that you should not write negative fractions like this: `-.3`, instead write them as `-0.3`.

Without any images, Arithmetic will act like a simple calculator and print the resulting output number on the standard output like the first example above. If you really want such calculator operations on the command-line, AWK (GNU AWK is the most common implementation) is much faster, easier and much more powerful. For example, the numerical one-line example above can be done with the following command. In general AWK is a fantastic tool and GNU AWK has a wonderful manual (`https://www.gnu.org/software/gawk/manual/`). So if you often confront situations like this, or have to work with large text tables/catalogs, be sure to checkout AWK and simplify your life.

```
$ echo "" | awk '{print (10.32-3.84)^2.7}'
155.329
```

6.3 Convolve

On an image, convolution can be thought of as a process to blur or remove the contrast in an image. If you are already familiar with the concept and just want to run Convolve, you can jump to Section 6.3.4 [Convolution kernel], page 135, and Section 6.3.5 [Invoking Convolve], page 136, and skip the lengthy introduction on the basic definitions and concepts of convolution.

There are generally two methods to convolve an image. The first and more intuitive one is in the "spatial domain" or using the actual image pixel values, see Section 6.3.1 [Spatial domain convolution], page 118. The second method is when we manipulate the "frequency domain", or work on the magnitudes of the different frequencies that constitute the image, see Section 6.3.2 [Frequency domain and Fourier operations], page 120. Understanding convolution in the spatial domain is more intuitive and thus recommended if you are just starting to learn about convolution. However, getting a good grasp of the frequency domain is a little more involved and needs some concentration and some mathematical proofs. However, its reward is a faster operation and more importantly a very fundamental understanding of this very important operation.

Convolution of an image will generally result in blurring the image because it mixes pixel values. In other words, if the image has sharp differences in neighboring pixel values[5], those sharp differences will become smoother. This has very good consequences in detection of signal in noise for example. In an actual observed image, the variation in neighboring pixel values due to noise can be very high. But after convolution, those variations will decrease and we have a better hope in detecting the possible underlying signal. Another case where convolution is extensively used is in mock images and modeling in general, convolution can be used to simulate the effect of the atmosphere or the optical system on the mock profiles that we create, see Section 8.1.1.2 [Point Spread Function], page 196. Convolution is a very interesting and important topic in any form of signal analysis (including astronomical observations). So we have thoroughly[6] explained the concepts behind it in the following sub-sections.

6.3.1 Spatial domain convolution

The pixels in an input image represent different "spatial" positions, therefore when convolution is done only using the actual input pixel values, we name the process as being done in the "Spatial domain". In particular this is in contrast to the "frequency domain" that we will discuss later in Section 6.3.2 [Frequency domain and Fourier operations], page 120. In the spatial domain (and in realistic situations where the image and the convolution kernel don't extend to infinity), convolution is the process of changing the value of one pixel to the *weighted* average of all the pixels in its *neighborhood*.

The 'neighborhood' of each pixel (how many pixels in which direction) and the 'weight' function (how much each neighboring pixel should contribute depending on its position) are given through a second image which is known as a "kernel"[7].

6.3.1.1 Convolution process

In convolution, the kernel specifies the weight and positions of the neighbors of each pixel. To find the convolved value of a pixel, the central pixel of the kernel is placed on that pixel. The values of each overlapping pixel in the kernel and image are multiplied by each other and summed for all the kernel pixels. To have one pixel in the center, the sides of the convolution kernel have to be an odd number. This process effectively mixes the pixel values of each pixel with its neighbors, resulting in a blurred image compared to the sharper input image.

Formally, convolution is one kind of linear 'spatial filtering' in image processing texts. If we assume that the kernel has $2a+1$ and $2b+1$ pixels on each side, the convolved value of a pixel placed at x and y ($C_{x,y}$) can be calculated from the neighboring pixel values in the input image (I) and the kernel (K) from

$$C_{x,y}=\sum_{s=-a}^{a}\sum_{t=-b}^{b}K_{s,t}\times I_{x+s,y+t}.$$

[5] In astronomy, the only major time we confront such sharp borders in signal are cosmic rays. All other sources of signal in an image are already blurred by the atmosphere or the optics of the instrument.

[6] A mathematician will certainly consider this explanation is incomplete and inaccurate. However this text is written for an understanding on the operations that are done on a real (not complex, discrete and noisy) astronomical image, not any general form of abstract function

[7] Also known as filter, here we will use 'kernel'.

Any pixel coordinate that is outside of the image in the equation above will be considered to be zero. When the kernel is symmetric about its center the blurred image has the same orientation as the original image. However, if the kernel is not symmetric, the image will be affected in the opposite manner, this is a natural consequence of the definition of spatial filtering. In order to avoid this we can rotate the kernel about its center by 180 degrees so the convolved output can have the same original orientation. Technically speaking, only if the kernel is flipped the process is known *Convolution*. If it isn't it is known as *Correlation*.

To be a weighted average, the sum of the weights (the pixels in the kernel) have to be unity. This will have the consequence that the convolved image of an object and unconvolved object will have the same brightness (see Section 8.1.3 [Flux Brightness and magnitude], page 201), which is natural, because convolution should not eat up the object photons, it only disperses them.

6.3.1.2 Edges in the spatial domain

In purely 'linear' spatial filtering (convolution), there are problems on the edges of the input image. Here we will explain the problem in the spatial domain. For a discussion of this problem from the frequency domain perspective, see Section 6.3.2.10 [Edges in the frequency domain], page 134. The problem originates from the fact that on the edges, in practice[8], the sum of the weights we use on the actual image pixels is not unity. For example, as discussed above, a profile in the center of an image will have the same brightness before and after convolution. However, for partially imaged profile on the edge of the image, the brightness (sum of its pixel fluxes within the image, see Section 8.1.3 [Flux Brightness and magnitude], page 201) will not be equal, some of the flux is going to be 'eaten' by the edges.

If you ran `$ make check` on the source files of Gnuastro, you can see the this effect by comparing the `convolve_frequency.fits` with `convolve_spatial.fits` in the `./tests/` directory. In the spatial domain, by default, no assumption will be made about pixels outside of the image or any blank pixels in the image. The problem explained above will also occur on the sides of blank regions (see Section 6.1.3 [Blank pixels], page 100). The solution to this edge effect problem is only possible in the spatial domain. For pixels near the edge, we have to abandon the assumption that the sum of the kernel pixels is unity during the convolution process[9]. So taking W as the sum of the kernel pixels that overlapped with non-blank and in-image pixels, the equation in Section 6.3.1.1 [Convolution process], page 118, will become:

$$C_{x,y}=\frac{\sum_{s=-a}^{a}\sum_{t=-b}^{b}K_{s,t}\times I_{x+s,y+t}}{W}.$$

In this manner, objects which are near the edges of the image or blank pixels will also have the same brightness (within the image) before and after convolution. This correction is applied by default in Convolve when convolving in the spatial domain. To disable it, you can use the `--noedgecorrection` option. In the frequency domain, there is no way to avoid

[8] Because we assumed the overlapping pixels outside the input image have a value of zero.

[9] ofcourse the sum of the kernel pixels still have to be unity in general.

this loss of flux near the edges of the image, see Section 6.3.2.10 [Edges in the frequency domain], page 134, for an interpretation from the frequency domain perspective.

Note that the edge effect discussed here is different from the one in Section 8.1.2 [If convolving afterwards], page 200. In making mock images we want to simulate a real observation. In a real observation the images of the galaxies on the sides of the CCD are first blurred by the atmosphere and instrument, then imaged. So light from the parts of a galaxy which are immediately outside the CCD will affect the parts of the galaxy which are covered by the CCD. Therefore in modeling the observation, we have to convolve an image that is larger than the input image by exactly half of the convolution kernel. We can hence conclude that this correction for the edges is only useful when working on actual observed images (where we don't have any more data on the edges) and not in modeling.

6.3.2 Frequency domain and Fourier operations

Getting a good grip on the frequency domain is usually not an easy job! So we have decided to give the issue a complete review here. Convolution in the frequency domain (see Section 6.3.2.6 [Convolution theorem], page 127) heavily relies on the concepts of Fourier transform (Section 6.3.2.4 [Fourier transform], page 125) and Fourier series (Section 6.3.2.3 [Fourier series], page 123) so we will be investigating these important operations first. It has become something of a cliché for people to say that the Fourier series "is a way to represent a (wave-like) function as the sum of simple sine waves" (from Wikipedia). However, sines themselves are abstract functions, so this statement really adds no extra layer of physical insight.

Before jumping head-first into the equations and proofs, we will begin with a historical background to see how the importance of frequencies actually roots in our ancient desire to see everything in terms of circles. A short review of how the complex plane should be interpreted is then given. Having paved the way with these two basics, we define the Fourier series and subsequently the Fourier transform. The final aim is to explain discrete Fourier transform, however some very important concepts need to be solidified first: The Dirac comb, convolution theorem and sampling theorem. So each of these topics are explained in their own separate sub-sub-section before going on to the discrete Fourier transform. Finally we revisit (after Section 6.3.1.2 [Edges in the spatial domain], page 119) the problem of convolution on the edges, but this time in the frequency domain. Understanding the sampling theorem and the discrete Fourier transform is very important in order to be able to pull out valuable science from the discrete image pixels. Therefore we have included the mathematical proofs and figures so you can have a clear understanding of these very important concepts.

6.3.2.1 Fourier series historical background

Ever since the ancient times, the circle has been (and still is) the simplest shape for abstract comprehension. All you need is a center point and a radius and you are done. All the points on a circle are at a fixed distance from the center. However, the moment you try to connect this elegantly simple and beautiful abstract construct (the circle) with the real world (for example compute its area or its circumference), things become really hard (ideally, impossible) because the irrational number π gets involved.

The key to understanding the Fourier series (thus the Fourier transform and finally the Discrete Fourier Transform) is our ancient desire to express everything in terms of circles or

the most exceptionally simple and elegant abstract human construct. Most people prefer to say the same thing in a more ahistorical manner: to break a function into sines and cosines. As the term "ancient" in the previous sentence implies, Jean-Baptiste Joseph Fourier (1768 – 1830 A.D.) was not the first person to do this. The main reason we know this process by his name today is that he came up with an ingenious method to find the necessary coefficients (radius of) and frequencies ("speed" of rotation on) the circles for any generic (integrable) function.

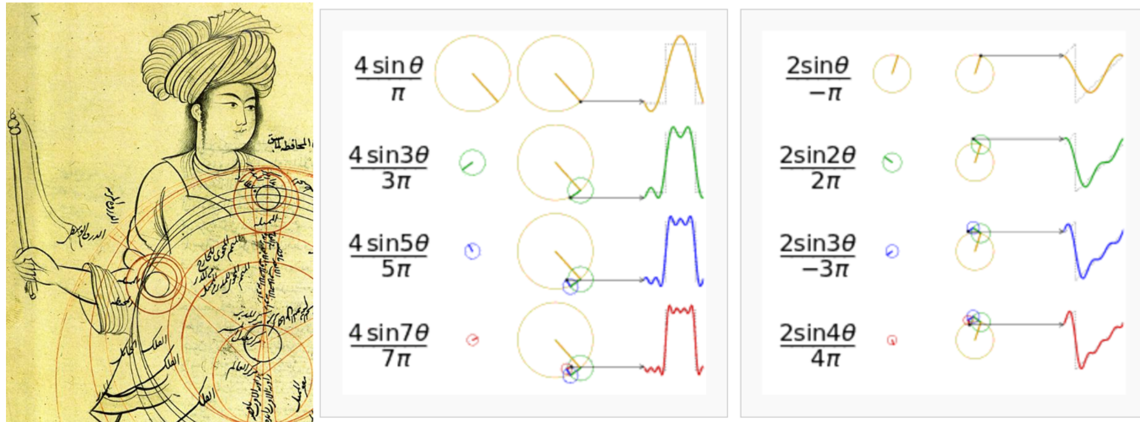

Figure 6.1: Epicycles and the Fourier series. Left: A demonstration of Mercury's epicycles relative to the "center of the world" by Qutb al-Din al-Shirazi (1236 – 1311 A.D.) retrieved from Wikipedia (`https://commons.wikimedia.org/wiki/File:Ghotb2.jpg`). Middle (`https://commons.wikimedia.org/wiki/File:Fourier_series_square_wave_circles_animation.gif`) and Right: How adding more epicycles (or terms in the Fourier series) will approximate functions. The right (`https://commons.wikimedia.org/wiki/File:Fourier_series_sawtooth_wave_circles_animation.gif`) animation is also available.

Like most aspects of mathematics, this process of interpreting everything in terms of circles, began for astronomical purposes. When astronomers noticed that the orbit of Mars and other outer planets, did not appear to be a simple circle (as everything should have been in the heavens). At some point during their orbit, the revolution of these planets would become slower, stop, go back a little (in what is known as the retrograde motion) and then continue going forward again.

The correction proposed by Ptolemy (90 – 168 A.D.) was the most agreed upon. He put the planets on Epicycles or circles whose center itself rotates on a circle whose center is the earth. Eventually, as observations became more and more precise, it was necessary to add more and more epicycles in order to explain the complex motions of the planets[10]. Figure 6.1(Left) shows an example depiction of the epicycles of Mercury in the late 13th century.

Of course we now know that if they had abdicated the Earth from its throne in the center of the heavens and allowed the Sun to take its place, everything would become much simpler

[10] See the Wikipedia page on "Deferent and epicycle" for a more complete historical review.

and true. But there wasn't enough observational evidence for changing the "professional consensus" of the time to this radical view suggested by a small minority[11]. So the pre-Galilean astronomers chose to keep Earth in the center and find a correction to the models (while keeping the heavens a purely "circular" order).

The main reason we are giving this historical background which might appear off topic is to give historical evidence that while such "approximations" do work and are very useful for pragmatic reasons (like measuring the calendar from the movement of astronomical bodies). They offer no physical insight. The astronomers who were involved with the Ptolemaic world view had to add a huge number of epicycles during the centuries after Ptolemy in order to explain more accurate observations. Finally the death knell of this world-view was Galileo's observations with his new instrument (the telescope). So the physical insight, which is what Astronomers and Physicists are interested in (as opposed to Mathematicians and Engineers who just like proving and optimizing or calculating!) comes from being creative and not limiting our selves to such approximations. Even when they work.

6.3.2.2 Circles and the complex plane

Before going onto the derivation, it is also useful to review how the complex numbers and their plane relate to the circles we talked about above. The two schematics in the middle and right of Figure 6.1 show how a 1D function of time can be made using the 2D real and imaginary surface. Seeing the animation in Wikipedia will really help in understanding this important concept. At each point in time, we take the vertical coordinate of the point and use it to find the value of the function at that point in time. Figure 6.2 shows this relation with the axes marked.

Leonhard Euler[12] (1707 – 1783 A.D.) showed that the complex exponential (e^{iv} where v is real) is periodic and can be written as: $e^{iv} = \cos v + isinv$. Therefore $e^{iv+2\pi} = e^{iv}$. Later, Caspar Wessel (mathematician and cartographer 1745 – 1818 A.D.) showed how complex numbers can be displayed as vectors on a plane. Euler's identity might seem counter intuitive at first, so we will try to explain it geometrically (for deeper physical insight). On the real-imaginary 2D plane (like the left hand plot in each box of Figure 6.2), multiplying a number by i can be interpretted as rotating the point by 90 degrees (for example the value 3 on the real axis becomes $3i$ on the imaginary axis). On the other hand, $e \equiv \lim_{n\rightarrow\infty}(1+\frac{1}{n})^n$, therefore, defining $m \equiv nu$, we get:

$$e^u = \lim_{n\rightarrow\infty}\left(1+\frac{1}{n}\right)^{nu} = \lim_{n\rightarrow\infty}\left(1+\frac{u}{nu}\right)^{nu} = \lim_{m\rightarrow\infty}\left(1+\frac{u}{m}\right)^{m}$$

Taking $u \equiv iv$ the result can be written as a generic complex number (a function of v):

---

[11] Aristarchus of Samos (310 – 230 B.C.) appears to be one of the first people to suggest the Sun being in the center of the universe. This approach to science (that the standard model is defined by consensus) and the fact that this consensus might be completely wrong still applies equally well to our models of particle physics and cosmology today.

[12] Other forms of this equation were known before Euler. For example in 1707 A.D. (the year of Euler's birth) Abraham de Moivre (1667 – 1754 A.D.) showed that $(\cos x + i\sin x)^n = \cos(nx) + i\sin(nx)$. In 1714 A.D., Roger Cotes (1682 – 1716 A.D. a colleague of Newton who proofread the second edition of Principia) showed that: $ix = \ln(\cos x + i\sin x)$.

$$e^{iv} = \lim_{m\to\infty}\left(1 + i\frac{v}{m}\right)^m = a(v) + ib(v)$$

For $v = \pi$, a nice geometric animation of going to the limit can be seen on Wikipedia (`https://commons.wikimedia.org/wiki/File:ExpIPi.gif`). We see that $\lim_{m\to\infty} a(\pi) = -1$, while $\lim_{m\to\infty} b(\pi) = 0$, which gives the famous $e^{i\pi} = -1$ equation. The final value is the real number -1, however the distance of the polygon points traversed as $m \to \infty$ is half the circumference of a circle or π, showing how v in the equation above can be interpretted as an angle in units of radians and therefore how $a(v) = cos(v)$ and $b(v) = sin(v)$.

Since e^{iv} is periodic (let's assume with a period of T), it is more clear to write it as $v \equiv \frac{2\pi n}{T}t$ (where n is an integer), so $e^{iv} = e^{i\frac{2\pi n}{T}t}$. The advantage of this notation is that the period (T) is clearly visible and the frequency ($\frac{2\pi n}{T}$, in units of 1/cycle) is defined through the integer n. In this notation, t is in units of "cycle"s.

As we see from the examples in Figure 6.1 and Figure 6.2, for each constituting frequency, we need a respective 'magnitude' or the radius of the circle in order to accurately approximate the desired 1D function. The concepts of "period" and "frequency" are relatively easy to grasp when using temporal units like time because this is how we define them in every-day life. However, in an image (astronomical data), we are dealing with spatial units like distance. Therefore, by one "period" we mean the *distance* at which the signal is identical and frequency is defined as the inverse of that spatial "period". The complex circle of Figure 6.2 can be thought of the Moon rotating about Earth which is rotating around the Sun; so the "Real (signal)" axis shows the Moon's position as seen by a distant observer on the Sun as time goes by. Because of the scalar (not having any direction or vector) nature of time, Figure 6.2 is easier to understand in units of time. When thinking about spatial units, mentally replace the "Time (sec)" axis with "Distance (meters)". Because length has direction and is a vector, visualizing the rotation of the imaginary circle and the advance along the "Distance (meters)" axis is not as simple as temporal units like time.

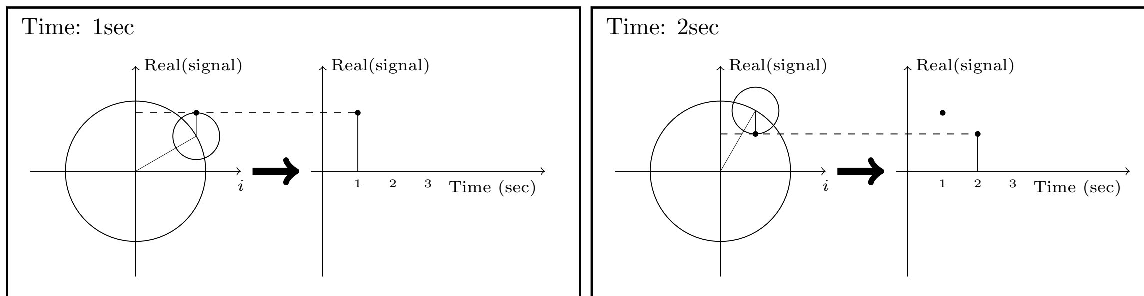

Figure 6.2: Relation between the real (signal), imaginary ($i \equiv \sqrt{-1}$) and time axes at two snapshots of time.

6.3.2.3 Fourier series

In astronomical images, our variable (brightness, or number of photo-electrons, or signal to be more generic) is recorded over the 2D spatial surface of a camera pixel. However to make things easier to understand, here we will assume that the signal is recorded in 1D (assume

one row of the 2D image pixels). Also for this section and the next (Section 6.3.2.4 [Fourier transform], page 125) we will be talking about the signal before it is digitized or pixelated. Let's assume that we have the continuous function $f(l)$ which is integrable in the interval $[l_0, l_0 + L]$ (always true in practical cases like images). Take l_0 as the position of the first pixel in the assumed row of the image and L as the width of the image along that row. The units of l_0 and L can be in any spatial units (for example meters) or an angular unit (like radians) multiplied by a fixed distance which is more common.

To approximate $f(l)$ over this interval, we need to find a set of frequencies and their corresponding 'magnitude's (see Section 6.3.2.2 [Circles and the complex plane], page 122). Therefore our aim is to show $f(l)$ as the following sum of periodic functions:

$$f(l) = \sum_{n=-\infty}^{\infty} c_n e^{i\frac{2\pi n}{L}l}$$

Note that the different frequencies ($2\pi n/L$, in units of cycles per meters for example) are not arbitrary. They are all integer multiples of the fundamental frequency of $\omega_0 = 2\pi/L$. Recall that L was the length of the signal we want to model. Therefore, we see that the smallest possible frequency (or the frequency resolution) in the end, depends on the length we observed the signal or L. In the case of each dimension on an image, this is the size of the image in the respective dimension. The frequencies have been defined in this "harmonic" fashion to insure that the final sum is periodic outside of the $[l_0, l_0 + L]$ interval too. At this point, you might be thinking that the sky is not periodic with the same period as my camera's view angle. You are absolutely right! The important thing is that since your camera's observed region is the only region we are "observing" and will be using, the rest of the sky is irrelevant; so we can safely assume the sky is periodic outside of it. However, this working assumption will haunt us later in Section 6.3.2.10 [Edges in the frequency domain], page 134.

The frequencies are thus determined by definition. So all we need to do is to find the coefficients (c_n), or magnitudes, or radii of the circles for each frequency which is identified with the integer n. Fourier's approach was to multiply both sides with a fixed term:

$$f(l)e^{-i\frac{2\pi m}{L}l} = \sum_{n=-\infty}^{\infty} c_n e^{i\frac{2\pi(n-m)}{L}l}$$

where $m > 0$[13]. We can then integrate both sides over the observation period:

$$\int_{l_0}^{l_0+L} f(l)e^{-i\frac{2\pi m}{L}l}dl = \int_{l_0}^{l_0+L} \sum_{n=-\infty}^{\infty} c_n e^{i\frac{2\pi(n-m)}{L}l}dl = \sum_{n=-\infty}^{\infty} c_n \int_{l_0}^{l_0+L} e^{i\frac{2\pi(n-m)}{L}l}dl$$

Both n and m are positive integers. Also, we know that a complex exponential is periodic so after one period (L) it comes back to its starting point. Therefore $\int_{l_0}^{l_0+L} e^{2\pi k/L}dl = 0$ for any $k > 0$. However, when $k = 0$, this integral becomes: $\int_{l_0}^{l_0+T} e^0 dt = \int_{l_0}^{l_0+T} dt = T$. Hence since the integral will be zero for all $n \neq m$, we get:

[13] We could have assumed $m < 0$ and set the exponential to positive, but this is more clear.

$$\sum_{n=-\infty}^{\infty} c_n \int_{l_0}^{l_0+T} e^{i\frac{2\pi(n-m)}{L}l}dl = Lc_m$$

The origin of the axis is fundamentally an arbitrary position. So let's set it to the start of the image such that $l_0 = 0$. So we can find the "magnitude" of the frequency $2\pi m/L$ within $f(l)$ through the relation:

$$c_m = \frac{1}{L}\int_0^L f(l)e^{-i\frac{2\pi m}{L}l}dl$$

6.3.2.4 Fourier transform

In Section 6.3.2.3 [Fourier series], page 123, we had to assume that the function is periodic outside of the desired interval with a period of L. Therefore, assuming that $L \to \infty$ will allow us to work with any function. However, with this approximation, the fundamental frequency (ω_0) or the frequency resolution that we discussed in Section 6.3.2.3 [Fourier series], page 123, will tend to zero: $\omega_0 \to 0$. In the equation to find c_m, every m represented a frequency (multiple of ω_0) and the integration on l removes the dependence of the right side of the equation on l, making it only a function of m or frequency. Let's define the following two variables:

$$\omega \equiv m\omega_0 = \frac{2\pi m}{L}$$

$$F(\omega) \equiv Lc_m$$

The equation to find the coefficients of each frequency in Section 6.3.2.3 [Fourier series], page 123, thus becomes:

$$F(\omega) = \int_{-\infty}^{\infty} f(l)e^{-i\omega l}dl.$$

The function $F(\omega)$ is thus the *Fourier transform* of $f(l)$ in the frequency domain. So through this transformation, we can find (analyze) the magnitudes of the constituting frequencies or the value in the frequency space[14] of our spatial input function. The great thing is that we can also do the reverse and later synthesize the input function from its Fourier transform. Let's do it: with the approximations above, multiply the right side of the definition of the Fourier Series (Section 6.3.2.3 [Fourier series], page 123) with $1 = L/L = (\omega_0 L)/(2\pi)$:

$$f(l) = \frac{1}{2\pi}\sum_{n=-\infty}^{\infty} Lc_n e^{\frac{2\pi in}{L}l}\omega_0 = \frac{1}{2\pi}\sum_{n=-\infty}^{\infty} F(\omega)e^{i\omega l}\Delta\omega$$

[14] As we discussed before, this 'magnitude' can be interpreted as the radius of the circle rotating at this frequency in the epicyclic interpretation of the Fourier series, see Figure 6.1 and Figure 6.2.

To find the right most side of this equation, we renamed ω_0 as $\Delta\omega$ because it was our resolution, $2\pi n/L$ was written as ω and finally, Lc_n was written as $F(\omega)$ as we defined above. Now, as $L \to \infty$, $\Delta\omega \to 0$ so we can write:

$$f(l) = \frac{1}{2\pi}\int_{-\infty}^{\infty} F(\omega)e^{i\omega l}d\omega$$

Together, these two equations provide us with a very powerful set of tools that we can use to process (analyze) and recreate (synthesize) the input signal. Through the first equation, we can break up our input function into its constituent frequencies and analyze it, hence it is also known as *analysis*. Using the second equation, we can synthesize or make the input function from the known frequencies and their magnitudes. Thus it is known as *synthesis*. Here, we symbolize the Fourier transform (analysis) and its inverse (synthesis) of a function $f(l)$ and its Fourier Transform $F(\omega)$ as $\mathcal{F}[f]$ and $\mathcal{F}^{-1}[F]$.

6.3.2.5 Dirac delta and comb

The Dirac δ (delta) function (also known as an impulse) is the way that we convert a continuous function into a discrete one. It is defined to satisfy the following integral:

$$\int_{-\infty}^{\infty} \delta(l)dl = 1$$

When integrated with another function, it gives that function's value at $l = 0$:

$$\int_{-\infty}^{\infty} f(l)\delta(l)dt = f(0)$$

An impulse positioned at another point (say l_0) is written as $\delta(l - l_0)$:

$$\int_{-\infty}^{\infty} f(l)\delta(l - l_0)dt = f(l_0)$$

The Dirac δ function also operates similarly if we use summations instead of integrals. The Fourier transform of the delta function is:

$$\mathcal{F}[\delta(l)] = \int_{-\infty}^{\infty} \delta(l)e^{-i\omega l}dl = e^{-i\omega 0} = 1$$

$$\mathcal{F}[\delta(l - l_0)] = \int_{-\infty}^{\infty} \delta(l - l_0)e^{-i\omega l}dl = e^{-i\omega l_0}$$

From the definition of the Dirac δ we can also define a Dirac comb (III_P) or an impulse train with infinite impulses separated by P:

$$\mathrm{III}_P(l) \equiv \sum_{k=-\infty}^{\infty} \delta(l - kP)$$

P is chosen to represent "pixel width" later in Section 6.3.2.7 [Sampling theorem], page 129. Therefore the Dirac comb is periodic with a period of P. We have intentionally used a different name for the period of the Dirac comb compared to the input signal's length of observation that we showed with L in Section 6.3.2.3 [Fourier series], page 123. This difference is highlighted here to avoid confusion later when these two periods are needed together in Section 6.3.2.8 [Discrete Fourier transform], page 132. The Fourier transform of the Dirac comb will be necessary in Section 6.3.2.7 [Sampling theorem], page 129, so let's derive it. By its definition, it is periodic, with a period of P, so the Fourier coefficients of its Fourier Series (Section 6.3.2.3 [Fourier series], page 123) can be calculated within one period:

$$\mathrm{III}_P = \sum_{n=-\infty}^{\infty} c_n e^{i\frac{2\pi n}{P}l}$$

We can now find the c_n from Section 6.3.2.3 [Fourier series], page 123:

$$c_n = \frac{1}{P}\int_{-P/2}^{P/2} \delta(l) e^{-i\frac{2\pi n}{P}l} = \frac{1}{P} \qquad \rightarrow \qquad \mathrm{III}_P = \frac{1}{P}\sum_{n=-\infty}^{\infty} e^{i\frac{2\pi n}{P}l}$$

So we can write the Fourier transform of the Dirac comb as:

$$\mathcal{F}[\mathrm{III}_P] = \int_{-\infty}^{\infty} \mathrm{III}_P e^{-i\omega l} dl = \frac{1}{P}\sum_{n=-\infty}^{\infty}\int_{-\infty}^{\infty} e^{-i(\omega - \frac{2\pi n}{P})l} dl = \frac{1}{P}\sum_{n=-\infty}^{\infty} \delta\left(\omega - \frac{2\pi n}{P}\right)$$

In the last step, we used the fact that the complex exponential is a periodic function, that n is an integer and that as we defined in Section 6.3.2.4 [Fourier transform], page 125, $\omega \equiv m\omega_0$, where m was an integer. The integral will be zero for any ω that is not equal to $2\pi n/P$, a more complete explanation can be seen in Section 6.3.2.3 [Fourier series], page 123. Therefore, while in the spatial domain the impulses had spacing of P (meters for example), in the frequency space, the spacing between the different impulses are $2\pi/P$ cycles per meters.

6.3.2.6 Convolution theorem

The convolution (shown with the $*$ operator) of the two functions $f(l)$ and $h(l)$ is defined as:

$$c(l) \equiv [f{*}h](l) = \int_{-\infty}^{\infty} f(\tau) h(l - \tau) d\tau$$

See Section 6.3.1.1 [Convolution process], page 118, for a more detailed physical (pixel based) interpretation of this definition. The Fourier transform of convolution ($C(\omega)$) can be written as:

$$C(\omega) = \int_{-\infty}^{\infty}[f{*}h](l)e^{-i\omega l}dl = \int_{-\infty}^{\infty}f(\tau)\left[\int_{-\infty}^{\infty}h(l-\tau)e^{-i\omega l}dl\right]d\tau$$

To solve the inner integral, let's define $s{\equiv}l-\tau$, so that $ds = dl$ and $l = s+\tau$ then the inner integral becomes:

$$\int_{-\infty}^{\infty}h(l-\tau)e^{-i\omega l}dl = \int_{-\infty}^{\infty}h(s)e^{-i\omega(s+\tau)}ds = e^{-i\omega\tau}\int_{-\infty}^{\infty}h(s)e^{-i\omega s}ds = H(\omega)e^{-i\omega\tau}$$

where $H(\omega)$ is the Fourier transform of $h(l)$. Substituting this result for the inner integral above, we get:

$$C(\omega) = H(\omega)\int_{-\infty}^{\infty}f(\tau)e^{-i\omega\tau}d\tau = H(\omega)F(\omega) = F(\omega)H(\omega)$$

where $F(\omega)$ is the Fourier transform of $f(l)$. So multiplying the Fourier transform of two functions individually, we get the Fourier transform of their convolution. The convolution theorem also proves a relation between the convolutions in the frequency space. Let's define:

$$D(\omega){\equiv}F(\omega){*}H(\omega)$$

Applying the inverse Fourier Transform or synthesis equation (Section 6.3.2.4 [Fourier transform], page 125) to both sides and following the same steps above, we get:

$$d(l) = f(l)h(l)$$

Where $d(l)$ is the inverse Fourier transform of $D(\omega)$. We can therefore re-write the two equations above formally as the convolution theorem:

$$\mathcal{F}[f{*}h] = \mathcal{F}[f]\mathcal{F}[h]$$

$$\mathcal{F}[fh] = \mathcal{F}[f] * \mathcal{F}[h]$$

Besides its usefulness in blurring an image by convolving it with a given kernel, the convolution theorem also enables us to do another very useful operation in data analysis: to match the blur (or PSF) between two images taken with different telescopes/cameras or under different atmospheric conditions. This process is also known as de-convolution. Let's

take $f(l)$ as the image with a narrower PSF (less blurry) and $c(l)$ as the image with a wider PSF which appears more blurred. Also let's take $h(l)$ to represent the kernel that should be convolved with the sharper image to create the more blurry image. Above, we proved the relation between these three images through the convolution theorem. But there, we assumed that $f(l)$ and $h(l)$ are known (given) and the convolved image is desired.

In de-convolution, we have $f(l)$ –the sharper image– and $f*h(l)$ –the more blurry image– and we want to find the kernel $h(l)$. The solution is a direct result of the convolution theorem:

$$\mathcal{F}[h]=\frac{\mathcal{F}[f*h]}{\mathcal{F}[f]} \qquad \text{or} \qquad h(l)=\mathcal{F}^{-1}\left[\frac{\mathcal{F}[f*h]}{\mathcal{F}[f]}\right]$$

While this works really nice, it has two problems:

- If $\mathcal{F}[f]$ has any zero values, then the inverse Fourier transform will not be a number!
- If there is significant noise in the image, then the high frequencies of the noise are going to significantly reduce the quality of the final result.

A standard solution to both these problems is the Weiner de-convolution algorithm[15].

6.3.2.7 Sampling theorem

Our mathematical functions are continuous, however, our data collecting and measuring tools are discrete. Here we want to give a mathematical formulation for digitizing the continuous mathematical functions so that later, we can retrieve the continuous function from the digitized recorded input. Assuming that we have a continuous function $f(l)$, then we can define $f_s(l)$ as the 'sampled' $f(l)$ through the Dirac comb (see Section 6.3.2.5 [Dirac delta and comb], page 126):

$$f_s(l)=f(l)\text{III}_P=\sum_{n=-\infty}^{\infty}f(l)\delta(l-nP)$$

The discrete data-element f_k (for example, a pixel in an image), where k is an integer, can thus be represented as:

$$f_k=\int_{-\infty}^{\infty}f_s(l)dl=\int_{-\infty}^{\infty}f(l)\delta(l-kP)dt=f(kP)$$

Note that in practice, our discrete data points are not found in this fashion. Each detector pixel (in an image for example) has an area and averages the signal it receives over that area, not a mathematical point as the Dirac δ function defines. However, as long as the variation in the signal over one detector pixel is not significant, this can be a good approximation. Having put this issue to the side, we can now try to find the relation between the Fourier transforms of the un-sampled $f(l)$ and the sampled $f_s(l)$. For a more clear notation, let's define:

[15] `https://en.wikipedia.org/wiki/Wiener_deconvolution`

$$F_s(\omega) \equiv \mathcal{F}[f_s]$$

$$D(\omega) \equiv \mathcal{F}[\mathrm{III}_P]$$

Then using the Convolution theorem (see Section 6.3.2.6 [Convolution theorem], page 127), $F_s(\omega)$ can be written as:

$$F_s(\omega) = \mathcal{F}[f(l)\mathrm{III}_P] = F(\omega)*D(\omega)$$

Finally, from the definition of convolution and the Fourier transform of the Dirac comb (see Section 6.3.2.5 [Dirac delta and comb], page 126), we get:

$$\begin{aligned} F_s(\omega) &= \int_{-\infty}^{\infty} F(\omega)D(\omega-\mu)d\mu \\ &= \frac{1}{P}\sum_{n=-\infty}^{\infty}\int_{-\infty}^{\infty} F(\omega)\delta\left(\omega-\mu-\frac{2\pi n}{P}\right)d\mu \\ &= \frac{1}{P}\sum_{n=-\infty}^{\infty} F\left(\omega-\frac{2\pi n}{P}\right). \end{aligned}$$

$F(\omega)$ was only a simple function, see Figure 6.3(left). However, from the sampled Fourier transform function we see that $F_s(\omega)$ is the superposition of infinite copies of $F(\omega)$ that have been shifted, see Figure 6.3(right). From the equation, it is clear that the shift in each copy is $2\pi/P$.

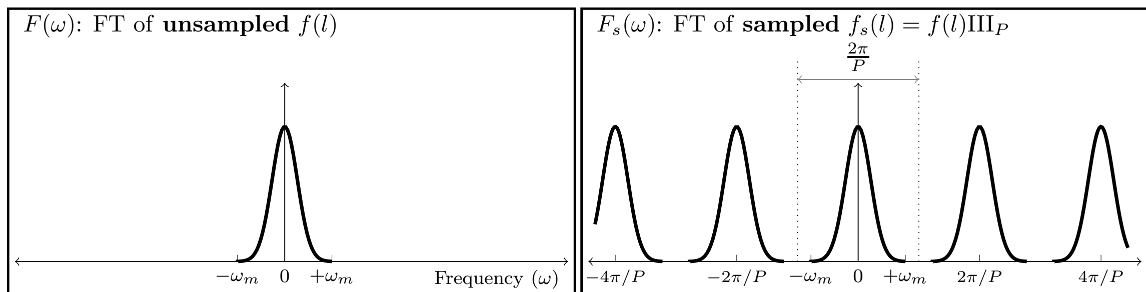

Figure 6.3: Sampling causes infinite repetition in the frequency domain. FT is an abbreviation for 'Fourier transform'. ω_m represents the maximum frequency present in the input. $F(\omega)$ is only symmetric on both sides of 0 when the input is real (not complex). In general $F(\omega)$ is complex and thus cannot be simply plotted like this. Here we have assumed a real Gaussian $f(t)$ which has produced a Gaussian $F(\omega)$.

The input $f(l)$ can have any distribution of frequencies in it. In the example of Figure 6.3(left), the input consisted of a range of frequencies equal to $\Delta\omega = 2\omega_m$.

Fortunately as Figure 6.3(right) shows, the assumed pixel size (P) we used to sample this hypothetical function was such that $2\pi/P > \Delta\omega$. The consequence is that each copy of $F(\omega)$ has become completely separate from the surrounding copies. Such a digitized (sampled) data set is thus called *over-sampled*. When $2\pi/P = \Delta\omega$, P is just small enough to finely separate even the largest frequencies in the input signal and thus it is known as *critically-sampled*. Finally if $2\pi/P < \Delta\omega$ we are dealing with an *under-sampled* data set. In an under-sampled data set, the separate copies of $F(\omega)$ are going to overlap and this will deprive us of recovering high constituent frequencies of $f(l)$. The effects of under-sampling in an image with high rates of change (for example a brick wall imaged from a distance) can clearly be visually seen and is known as *aliasing*.

When the input $f(l)$ is composed of a finite range of frequencies, $f(l)$ is known as a *band-limited* function. The example in Figure 6.3(left) was a nice demonstration of such a case: for all $\omega < -\omega_m$ or $\omega > \omega_m$, we have $F(\omega) = 0$. Therefore, when the input function is band-limited and our detector's pixels are placed such that we have critically (or over-) sampled it, then we can exactly reproduce the continuous $f(l)$ from the discrete or digitized samples. To do that, we just have to isolate one copy of $F(\omega)$ from the infinite copies and take its inverse Fourier transform.

This ability to exactly reproduce the continuous input from the sampled or digitized data leads us to the *sampling theorem* which connects the inherent property of the continuous signal (its maximum frequency) to that of the detector (the spacing between its pixels). The sampling theorem states that the full (continuous) signal can be recovered when the pixel size (P) and the maximum constituent frequency in the signal (ω_m) have the following relation[16]:

$$\frac{2\pi}{P} > 2\omega_m$$

This relation was first formulated by Harry Nyquist (1889 – 1976 A.D.) in 1928 and formally proved in 1949 by Claude E. Shannon (1916 – 2001 A.D.) in what is now known as the Nyquist-Shannon sampling theorem. In signal processing, the signal is produced (synthesized) by a transmitter and is received and de-coded (analyzed) by a receiver. Therefore producing a band-limited signal is necessary.

In astronomy, we do not produce the shapes of our targets, we are only observers. Galaxies can have any shape and size, therefore ideally, our signal is not band-limited. However, since we are always confined to observing through an aperture, the aperture will cause a point source (for which $\omega_m = \infty$) to be spread over several pixels. This spread is quantitatively known as the point spread function or PSF. This spread does blur the image which is undesirable; however, for this analysis it produces the positive outcome that there will be a finite ω_m. Though we should caution that any detector will have noise which will add lots of very high frequency (ideally infinite) changes between the pixels. However, the coefficients of those noise frequencies are usually exceedingly small.

[16] This equation is also shown in some places without the 2π. Whether 2π is included or not depends on how you define the frequency

6.3.2.8 Discrete Fourier transform

As we have stated several times so far, the input image is a digitized, pixelated or discrete array of values ($f_s(l)$, see Section 6.3.2.7 [Sampling theorem], page 129). The input is not a continuous function. Also, all our numerical calculations can only be done on a sampled, or discrete Fourier transform. Note that $F_s(\omega)$ is not discrete, it is continuous. One way would be to find the analytic $F_s(\omega)$, then sample it at any desired "freq-pixel"[17] spacing. However, this process would involve two steps of operations and computers in particular are not too good at analytic operations for the first step. So here, we will derive a method to directly find the 'freq-pixel'ated $F_s(\omega)$ from the pixelated $f_s(l)$. Let's start with the definition of the Fourier transform (see Section 6.3.2.4 [Fourier transform], page 125):

$$F_s(\omega) = \int_{-\infty}^{\infty} f_s(l)e^{-i\omega l}dl$$

From the definition of $f_s(\omega)$ (using x instead of n) we get:

$$\begin{aligned} F_s(\omega) &= \sum_{x=-\infty}^{\infty} \int_{-\infty}^{\infty} f(l)\delta(l - xP)e^{-i\omega l}dl \\ &= \sum_{x=-\infty}^{\infty} f_x e^{-i\omega xP} \end{aligned}$$

Where f_x is the value of $f(l)$ on the point x or the value of the xth pixel. As shown in Section 6.3.2.7 [Sampling theorem], page 129, this function is infinitely periodic with a period of $2\pi/P$. So all we need is the values within one period: $0 < \omega < 2\pi/P$, see Figure 6.3. We want X samples within this interval, so the frequency difference between each frequency sample or freq-pixel is $1/XP$. Hence we will evaluate the equation above on the points at:

$$\omega = \frac{u}{XP} \qquad u = 0, 1, 2, ..., X - 1$$

Therefore the value of the freq-pixel u in the frequency domain is:

$$F_u = \sum_{x=0}^{X-1} f_x e^{-i\frac{ux}{X}}$$

Therefore, we see that for each freq-pixel in the frequency domain, we are going to need all the pixels in the spatial domain[18]. If the input (spatial) pixel row is also X pixels wide, then we can exactly recover the xth pixel with the following summation:

$$f_x = \frac{1}{X}\sum_{u=0}^{X-1} F_u e^{i\frac{ux}{X}}$$

[17] We are using the made-up word "freq-pixel" so they are not confused with spatial domain "pixels".

[18] So even if one pixel is a blank pixel (see Section 6.1.3 [Blank pixels], page 100), all the pixels in the frequency domain will also be blank.

When the input pixel row (we are still only working on 1D data) has X pixels, then it is $L = XP$ spatial units wide. L, or the length of the input data was defined in Section 6.3.2.3 [Fourier series], page 123, and P or the space between the pixels in the input was defined in Section 6.3.2.5 [Dirac delta and comb], page 126. As we saw in Section 6.3.2.7 [Sampling theorem], page 129, the input (spatial) pixel spacing (P) specifies the range of frequencies that can be studied and in Section 6.3.2.3 [Fourier series], page 123, we saw that the length of the (spatial) input, (L) determines the resolution (or size of the freq-pixels) in our discrete Fourier transformed image. Both result from the fact that the frequency domain is the inverse of the spatial domain.

6.3.2.9 Fourier operations in two dimensions

Once all the relations in the previous sections have been clearly understood in one dimension, it is very easy to generalize them to two or even more dimensions since each dimension is by definition independent. Previously we defined l as the continuous variable in 1D and the inverse of the period in its direction to be ω. Let's show the second spatial direction with m the the inverse of the period in the second dimension with ν. The Fourier transform in 2D (see Section 6.3.2.4 [Fourier transform], page 125) can be written as:

$$F(\omega,\nu)=\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}f(l,m)e^{-i(\omega l+\nu m)}dl$$

$$f(l,m)=\int_{-\infty}^{\infty}\int_{-\infty}^{\infty}F(\omega,\nu)e^{i(\omega l+\nu m)}dl$$

The 2D Dirac $\delta(l,m)$ is non-zero only when $l = m = 0$. The 2D Dirac comb (or Dirac brush! See Section 6.3.2.5 [Dirac delta and comb], page 126) can be written in units of the 2D Dirac δ. For most image detectors, the sides of a pixel are equal in both dimensions. So P remains unchanged, if a specific device is used which has non-square pixels, then for each dimension a different value should be used.

$$\mathrm{III}_P(l,m)\equiv\sum_{j=-\infty}^{\infty}\sum_{k=-\infty}^{\infty}\delta(l-jP,m-kP)$$

The Two dimensional Sampling theorem (see Section 6.3.2.7 [Sampling theorem], page 129) is thus very easily derived as before since the frequencies in each dimension are independent. Let's take ν_m as the maximum frequency along the second dimension. Therefore the two dimensional sampling theorem says that a 2D band-limited function can be recovered when the following conditions hold[19]:

$$\frac{2\pi}{P}>2\omega_m \qquad \text{and} \qquad \frac{2\pi}{P}>2\nu_m$$

[19] If the pixels are not a square, then each dimension has to use the respective pixel size, but since most imagers have square pixels, we assume so here too

Finally, let's represent the pixel counter on the second dimension in the spatial and frequency domains with y and v respectively. Also let's assume that the input image has Y pixels on the second dimension. Then the two dimensional discrete Fourier transform and its inverse (see Section 6.3.2.8 [Discrete Fourier transform], page 132) can be written as:

$$F_{u,v}=\sum_{x=0}^{X-1}\sum_{y=0}^{Y-1}f_{x,y}e^{-i(\frac{ux}{X}+\frac{vy}{Y})}$$

$$f_{x,y}=\frac{1}{XY}\sum_{u=0}^{X-1}\sum_{v=0}^{Y-1}F_{u,v}e^{i(\frac{ux}{X}+\frac{vy}{Y})}$$

6.3.2.10 Edges in the frequency domain

With a good grasp of the frequency domain, we can revisit the problem of convolution on the image edges, see Section 6.3.1.2 [Edges in the spatial domain], page 119. When we apply the convolution theorem (see Section 6.3.2.6 [Convolution theorem], page 127) to convolve an image, we first take the discrete Fourier transforms (DFT, Section 6.3.2.8 [Discrete Fourier transform], page 132) of both the input image and the kernel, then we multiply them with each other and then take the inverse DFT to construct the convolved image. Of course, in order to multiply them with each other in the frequency domain, the two images have to be the same size, so let's assume that we pad the kernel (it is usually smaller than the input image) with zero valued pixels in both dimensions so it becomes the same size as the input image before the DFT.

Having multiplied the two DFTs, we now apply the inverse DFT which is where the problem is usually created. If the DFT of the kernel only had values of 1 (unrealistic condition!) then there would be no problem and the inverse DFT of the multiplication would be identical with the input. However in real situations, the kernel's DFT has a maximum of 1 (because the sum of the kernel has to be one, see Section 6.3.1.1 [Convolution process], page 118) and decreases something like the hypothetical profile of Figure 6.3. So when multiplied with the input image's DFT, the coefficients or magnitudes (see Section 6.3.2.2 [Circles and the complex plane], page 122) of the smallest frequency (or the sum of the input image pixels) remains unchanged, while the magnitudes of the higher frequencies are significantly reduced.

As we saw in Section 6.3.2.7 [Sampling theorem], page 129, the Fourier transform of a discrete input will be infinitely repeated. In the final inverse DFT step, the input is in the frequency domain (the multiplied DFT of the input image and the kernel DFT). So the result (our output convolved image) will be infinitely repeated in the spatial domain. In order to accurately reconstruct the input image, we need all the frequencies with the correct magnitudes. However, when the magnitudes of higher frequencies are decreased, longer periods (shorter frequencies) will dominate in the reconstructed pixel values. Therefore, when constructing a pixel on the edge of the image, the newly empowered longer periods will look beyond the input image edges and will find the repeated input image there. So if you convolve an image in this fashion using the convolution theorem, when a bright object exists on one edge of the image, its blurred wings will be present on the other side of the convolved image. This is often termed as circular convolution or cyclic convolution.

So, as long as we are dealing with convolution in the frequency domain, there is nothing we can do about the image edges. The least we can do is to eliminate the ghosts of the other side of the image. So, we add zero valued pixels to both the input image and the kernel in both dimensions so the image that will be convolved has a size equal to the sum of both images in each dimension. Of course, the effect of this zero-padding is that the sides of the output convolved image will become dark. To put it another way, the edges are going to drain the flux from nearby objects. But at least it is consistent across all the edges of the image and is predictable. In Convolve, you can see the padded images when inspecting the frequency domain convolution steps with the `--viewfreqsteps` option.

6.3.3 Spatial vs. Frequency domain

With the discussions above it might not be clear when to choose the spatial domain and when to choose the frequency domain. Here we will try to list the benefits of each.

The spatial domain,

- Can correct for the edge effects of convolution, see Section 6.3.1.2 [Edges in the spatial domain], page 119.
- Can operate on blank pixels.
- Can be faster than frequency domain when the kernel is small (in terms of the number of pixels on the sides).

The frequency domain,

- Will be much faster when the image and kernel are both large.

As a general rule of thumb, when working on an image of modeled profiles use the frequency domain and when working on an image of real (observed) objects use the spatial domain (corrected for the edges). The reason is that if you apply a frequency domain convolution to a real image, you are going to loose information on the edges and generally you don't want large kernels. But when you have made the profiles in the image yourself, you can just make a larger input image and crop the central parts to completely remove the edge effect, see Section 8.1.2 [If convolving afterwards], page 200. Also due to oversampling, both the kernels and the images can become very large and the speed boost of frequency domain convolution will significantly improve the processing time, see Section 8.1.1.6 [Oversampling], page 200.

6.3.4 Convolution kernel

All the programs that need convolution will need to be given a convolution kernel file and extension. In most cases (other than Convolve, see Section 6.3 [Convolve], page 117) the kernel file name is optional. However, the extension is necessary and must be specified either on the command-line or at least one of the configuration files (see Section 4.2 [Configuration files], page 59). Within Gnuastro, there are two ways to create a kernel image:

- MakeProfiles: You can use MakeProfiles to create a parametric (based on a radial function) kernel, see Section 8.1 [MakeProfiles], page 195. By default MakeProfiles will make the Gaussian and Moffat profiles in a separate file so you can feed it into any of the programs.
- ConvertType: You can write your own desired kernel into a text file table and convert it to a FITS file with ConvertType, see Section 5.2 [ConvertType], page 87. Just be careful that the kernel has to have an odd number of pixels along its two axes, see

Section 6.3.1.1 [Convolution process], page 118. All the programs that do convolution will normalize the kernel internally, so if you choose this option, you don't have to worry about normalizing the kernel. Only within Convolve, there is an option to disable normalization, see Section 6.3.5 [Invoking Convolve], page 136.

The two options to specify a kernel file name and its extension are shown below. These are common between all the programs that will do convolution.

`-k STR`
`--kernel=STR`
: The convolution kernel file name. The `BITPIX` (data type) value of this file can be any standard type and it does not necessarily have to be normalized. Several operations will be done on the kernel image prior to the program's processing:

- It will be converted to floating point type.
- All blank pixels (see Section 6.1.3 [Blank pixels], page 100) will be set to zero.
- It will be normalized so the sum of its pixels equal unity.
- It will be flipped so the convolved image has the same orientation. This is only relevant if the kernel is not circular. See Section 6.3.1.1 [Convolution process], page 118.

`-U STR`
`--khdu=STR`
: The convolution kernel HDU. Although the kernel file name is optional, before running any of the programs, they need to have a value for `--khdu` even if the default kernel is to be used. So be sure to keep its value in at least one of the configuration files (see Section 4.2 [Configuration files], page 59). By default, the system configuration file has a value.

6.3.5 Invoking Convolve

Convolve an input image with a known kernel or make the kernel necessary to match two PSFs. The general template for Convolve is:

```
$ astconvolve [OPTION...] ASTRdata
```

One line examples:

```
## Convolve mockimg.fits with psf.fits:
$ astconvolve --kernel=psf.fits mockimg.fits

## Convolve in the spatial domain:
$ astconvolve observedimg.fits --kernel=psf.fits --domain=spatial

## Find the kernel to match sharper and blurry PSF images:
$ astconvolve --kernel=sharperimage.fits --makekernel=10           \
              blurryimage.fits
```

The only argument accepted by Convolve is an input image file. Some of the options are the same between Convolve and some other Gnuastro programs. Therefore, to avoid repetition, they will not be repeated here. For the full list of options shared by all Gnuastro programs, please see Section 4.1.2 [Common options], page 52. In particular, in the spatial

domain convolve uses Gnuastro's tessellation, see Section 4.6 [Tessellation], page 72, and the common options related to that in Section 4.1.2.2 [Processing options], page 54.

Here we will only explain the options particular to Convolve. Run Convolve with `--help` in order to see the full list of options Convolve accepts, irrespective of where they are explained in this book.

`--nokernelflip`
: Do not flip the kernel after reading it the spatial domain convolution. This can be useful if the flipping has already been applied to the kernel.

`--nokernelnorm`
: Do not normalize the kernel after reading it, such that the sum of its pixels is unity.

`-d STR`
`--domain=STR`
: The domain to use for the convolution. The acceptable values are '`spatial`' and '`frequency`', corresponding to the respective domain.

  For large images, the frequency domain process will be more efficient than convolving in the spatial domain. However, the edges of the image will loose some flux (see Section 6.3.1.2 [Edges in the spatial domain], page 119) and the image must not contain any blank pixels, see Section 6.3.3 [Spatial vs. Frequency domain], page 135.

`--checkfreqsteps`
: With this option a file with the initial name of the output file will be created that is suffixed with `_freqsteps.fits`, all the steps done to arrive at the final convolved image are saved as extensions in this file. The extensions in order are:

  1. The padded input image. In frequency domain convolution the two images (input and convolved) have to be the same size and both should be padded by zeros.
  2. The padded kernel, similar to the above.
  3. The Fourier spectrum of the forward Fourier transform of the input image. Note that the Fourier transform is a complex operation (and not view able in one image!) So we either have to show the 'Fourier spectrum' or the 'Phase angle'. For the complex number $a + ib$, the Fourier spectrum is defined as $\sqrt{a^2 + b^2}$ while the phase angle is defined as $\arctan(b/a)$.
  4. The Fourier spectrum of the forward Fourier transform of the kernel image.
  5. The Fourier spectrum of the multiplied (through complex arithmetic) transformed images.
  6. The inverse Fourier transform of the multiplied image. If you open it, you will see that the convolved image is now in the center, not on one side of the image as it started with (in the padded image of the first extension). If you are working on a mock image which originally had pixels of precisely 0.0, you will notice that in those parts that your convolved profile(s) did not convert, the values are now $\sim 10^{-18}$, this is due to floating-point round

off errors. Therefore in the final step (when cropping the central parts of the image), we also remove any pixel with a value less than 10^{-17}.

`--noedgecorrection`
Do not correct the edge effect in spatial domain convolution. For a full discussion, please see Section 6.3.1.2 [Edges in the spatial domain], page 119.

`-m INT`
`--makekernel=INT`
(`=INT`) If this option is called, Convolve will do de-convolution (see Section 6.3.2.6 [Convolution theorem], page 127). The image specified by the `--kernel` option is assumed to be the sharper (less blurry) image and the input image is assumed to be the more blurry image. The value given to this option will be used as the maximum radius of the kernel. Any pixel in the final kernel that is larger than this distance from the center will be set to zero. The two images must have the same size.

Noise has large frequencies which can make the result less reliable for the higher frequencies of the final result. So all the frequencies which have a spectrum smaller than the value given to the `minsharpspec` option in the sharper input image are set to zero and not divided. This will cause the wings of the final kernel to be flatter than they would ideally be which will make the convolved image result unreliable if it is too high. Some notes to take into account for a good result:

- Choose a bright (unsaturated) star and use a region box (with Crop for example, see Section 6.1 [Crop], page 97) that is sufficiently above the noise.
- Use Warp (see Section 6.4 [Warp], page 138) to warp the pixel grid so the star's center is exactly on the center of the central pixel in the cropped image. This will certainly slightly degrade the result, however, it is necessary. If there are multiple good stars, you can shift all of them, then normalize them (so the sum of each star's pixels is one) and then take their average to decrease this effect.
- The shifting might move the center of the star by one pixel in any direction, so crop the central pixel of the warped image to have a clean image for the de-convolution.

`-c`
`--minsharpspec`
(`=FLT`) The minimum frequency spectrum (or coefficient, or pixel value in the frequency domain image) to use in deconvolution, see the explanations under the `--makekernel` option for more information.

6.4 Warp

Image warping is the process of mapping the pixels of one image onto a new pixel grid. This process is sometimes known as transformation, however following the discussion of

Heckbert 1989[20] we will not be using that term because it can be confused with only pixel value or flux transformations. Here we specifically mean the pixel grid transformation which is better conveyed with 'warp'.

Image wrapping is a very important step in astronomy, both in observational data analysis and in simulating modeled images. In modeling, warping an image is necessary when we want to apply grid transformations to the initial models, for example in simulating gravitational lensing (Radial warpings are not yet included in Warp). Observational reasons for warping an image are listed below:

- **Noise:** Most scientifically interesting targets are inherently faint (have a very low Signal to noise ratio). Therefore one short exposure is not enough to detect such objects that are drowned deeply in the noise. We need multiple exposures so we can add them together and increase the objects' signal to noise ratio. Keeping the telescope fixed on one field of the sky is practically impossible. Therefore very deep observations have to put into the same grid before adding them.
- **Resolution:** If we have multiple images of one patch of the sky (hopefully at multiple orientations) we can warp them to the same grid. The multiple orientations will allow us to 'guess' the values of pixels on an output pixel grid that has smaller pixel sizes and thus increase the resolution of the output. This process of merging multiple observations is known as Mosaicing.
- **Cosmic rays:** Cosmic rays can randomly fall on any part of an image. If they collide vertically with the camera, they are going to create a very sharp and bright spot that in most cases can be separated easily[21]. However, depending on the depth of the camera pixels, and the angle that a cosmic rays collides with it, it can cover a line-like larger area on the CCD which makes the detection using their sharp edges very hard and error prone. One of the best methods to remove cosmic rays is to compare multiple images of the same field. To do that, we need all the images to be on the same pixel grid.
- **Optical distortion:** (Not yet included in Warp) In wide field images, the optical distortion that occurs on the outer parts of the focal plane will make accurate comparison of the objects at various locations impossible. It is therefore necessary to warp the image and correct for those distortions prior to the analysis.
- **Detector not on focal plane:** In some cases (like the Hubble Space Telescope ACS and WFC3 cameras), the CCD might be tilted compared to the focal plane, therefore the recorded CCD pixels have to be projected onto the focal plane before further analysis.

6.4.1 Warping basics

Let's take $[\,u \quad v\,]$ as the coordinates of a point in the input image and $[\,x \quad y\,]$ as the coordinates of that same point in the output image[22]. The simplest form of coordinate transformation (or warping) is the scaling of the coordinates, let's assume we want to scale the first axis by M and the second by N, the output coordinates of that point can be calculated by

[20] Paul S. Heckbert. 1989. *Fundamentals of Texture mapping and Image Warping*, Master's thesis at University of California, Berkely.

[21] All astronomical targets are blurred with the PSF, see Section 8.1.1.2 [Point Spread Function], page 196, however a cosmic ray is not and so it is very sharp (it suddenly stops at one pixel).

[22] These can be any real number, we are not necessarily talking about integer pixels here.

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} Mu \\ Nv \end{bmatrix} = \begin{bmatrix} M & 0 \\ 0 & N \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}$$

Note that these are matrix multiplications. We thus see that we can represent any such grid warping as a matrix. Another thing we can do with this 2×2 matrix is to rotate the output coordinate around the common center of both coordinates. If the output is rotated anticlockwise by θ degrees from the positive (to the right) horizontal axis, then the warping matrix should become:

$$\begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} ucos\theta - vsin\theta \\ usin\theta + vcos\theta \end{bmatrix} = \begin{bmatrix} cos\theta & -sin\theta \\ sin\theta & cos\theta \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}$$

We can also flip the coordinates around the first axis, the second axis and the coordinate center with the following three matrices respectively:

$$\begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \qquad \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix} \qquad \begin{bmatrix} -1 & 0 \\ 0 & -1 \end{bmatrix}$$

The final thing we can do with this definition of a 2×2 warping matrix is shear. If we want the output to be sheared along the first axis with A and along the second with B, then we can use the matrix:

$$\begin{bmatrix} 1 & A \\ B & 1 \end{bmatrix}$$

To have one matrix representing any combination of these steps, you use matrix multiplication, see Section 6.4.2 [Merging multiple warpings], page 142. So any combinations of these transformations can be displayed with one 2×2 matrix:

$$\begin{bmatrix} a & b \\ c & d \end{bmatrix}$$

The transformations above can cover a lot of the needs of most coordinate transformations. However they are limited to mapping the point $[0 \quad 0]$ to $[0 \quad 0]$. Therefore they are useless if you want one coordinate to be shifted compared to the other one. They are also space invariant, meaning that all the coordinates in the image will receive the same transformation. In other words, all the pixels in the output image will have the same area if placed over the input image. So transformations which require varying output pixel sizes like projections cannot be applied through this 2×2 matrix either (for example for the tilted ACS and WFC3 camera detectors on board the Hubble space telescope).

To add these further capabilities, namely translation and projection, we use the homogeneous coordinates. They were defined about 200 years ago by August Ferdinand Möbius (1790 – 1868). For simplicity, we will only discuss points on a 2D plane and avoid the complexities of higher dimensions. We cannot provide a deep mathematical introduction here,

interested readers can get a more detailed explanation from Wikipedia[23] and the references therein.

By adding an extra coordinate to a point we can add the flexibility we need. The point $[\,x \quad y\,]$ can be represented as $[\,xZ \quad yZ \quad Z\,]$ in homogeneous coordinates. Therefore multiplying all the coordinates of a point in the homogeneous coordinates with a constant will give the same point. Put another way, the point $[\,x \quad y \quad Z\,]$ corresponds to the point $[\,x/Z \quad y/Z\,]$ on the constant Z plane. Setting $Z = 1$, we get the input image plane, so $[\,u \quad v \quad 1\,]$ corresponds to $[\,u \quad v\,]$. With this definition, the transformations above can be generally written as:

$$\begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \begin{bmatrix} a & b & 0 \\ c & d & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} u \\ v \\ 1 \end{bmatrix}$$

We thus acquired 4 extra degrees of freedom. By giving non-zero values to the zero valued elements of the last column we can have translation (try the matrix multiplication!). In general, any coordinate transformation that is represented by the matrix below is known as an affine transformation[24]:

$$\begin{bmatrix} a & b & c \\ d & e & f \\ 0 & 0 & 1 \end{bmatrix}$$

We can now consider translation, but the affine transform is still spatially invariant. Giving non-zero values to the other two elements in the matrix above gives us the projective transformation or Homography[25] which is the most general type of transformation with the 3×3 matrix:

$$\begin{bmatrix} x' \\ y' \\ w \end{bmatrix} = \begin{bmatrix} a & b & c \\ d & e & f \\ g & h & 1 \end{bmatrix} \begin{bmatrix} u \\ v \\ 1 \end{bmatrix}$$

So the output coordinates can be calculated from:

$$x = \frac{x'}{w} = \frac{au + bv + c}{gu + hv + 1} \qquad\qquad y = \frac{y'}{w} = \frac{du + ev + f}{gu + hv + 1}$$

Thus with homography we can change the sizes of the output pixels on the input plane, giving a 'perspective'-like visual impression. This can be quantitatively seen in the two equations above. When $g = h = 0$, the denominator is independent of u or v and thus we have spatial invariance. Homography preserves lines at all orientations. A very useful fact about homography is that its inverse is also a homography. These two properties play a very important role in the implementation of this transformation. A short but instructive and

---

[23] `http://en.wikipedia.org/wiki/Homogeneous_coordinates`

[24] `http://en.wikipedia.org/wiki/Affine_transformation`

[25] `http://en.wikipedia.org/wiki/Homography`

illustrated review of affine, projective and also bi-linear mappings is provided in Heckbert 1989[26].

6.4.2 Merging multiple warpings

In Section 6.4.1 [Warping basics], page 139, we saw how one basic warping/transformation can be represented with a 3 by 3 matrix. To make more complex warpings these matrices have to be multiplied through matrix multiplication. However matrix multiplication is not commutative, so the order of the set of matrices you use for the multiplication is going to be very important.

The first warping should be placed as the left-most matrix. The second warping to the right of that and so on. The second transformation is going to occur on the warped coordinates of the first. As an example for merging a few transforms into one matrix, the multiplication below represents the rotation of an image about a point $[U \quad V]$ anticlockwise from the horizontal axis by an angle of θ. To do this, first we take the origin to $[U \quad V]$ through translation. Then we rotate the image, then we translate it back to where it was initially. These three operations can be merged in one operation by calculating the matrix multiplication below:

$$\begin{bmatrix}1 & 0 & U\\ 0 & 1 & V\\ 0 & 0 & 1\end{bmatrix}\begin{bmatrix}cos\theta & -sin\theta & 0\\ sin\theta & cos\theta & 0\\ 0 & 0 & 1\end{bmatrix}\begin{bmatrix}1 & 0 & -U\\ 0 & 1 & -V\\ 0 & 0 & 1\end{bmatrix}$$

6.4.3 Resampling

A digital image is composed of discrete 'picture elements' or 'pixels'. When a real image is created from a camera or detector, each pixel's area is used to store the number of photo-electrons that were created when incident photons collided with that pixel's surface area. This process is called the 'sampling' of a continuous or analog data into digital data. When we change the pixel grid of an image or warp it as we defined in Section 6.4.1 [Warping basics], page 139, we have to 'guess' the flux value of each pixel on the new grid based on the old grid, or re-sample it. Because of the 'guessing', any form of warping on the data is going to degrade the image and mix the original pixel values with each other. So if an analysis can be done on an un-warped data image, it is best to leave the image untouched and pursue the analysis. However as discussed in Section 6.4 [Warp], page 138, this is not possible most of the times, so we have to accept the problem and re-sample the image.

In most applications of image processing, it is sufficient to consider each pixel to be a point and not an area. This assumption can significantly speed up the processing of an image and also the simplicity of the code. It is a fine assumption when the signal to noise ratio of the objects are very large. The question will then be one of interpolation because you have multiple points distributed over the output image and you want to find the values at the pixel centers. To increase the accuracy, you might also sample more than one point from within a pixel giving you more points for a more accurate interpolation in the output grid.

---

[26] Paul S. Heckbert. 1989. *Fundamentals of Texture mapping and Image Warping*, Master's thesis at University of California, Berkely. Note that since points are defined as row vectors there, the matrix is the transpose of the one discussed here.

However, interpolation has several problems. The first one is that it will depend on the type of function you want to assume for the interpolation. For example you can choose a bi-linear or bi-cubic (the 'bi's are for the 2 dimensional nature of the data) interpolation method. For the latter there are various ways to set the constants[27]. Such functional interpolation functions can fail seriously on the edges of an image. They will also need normalization so that the flux of the objects before and after the warpings are comparable. The most basic problem with such techniques is that they are based on a point while a detector pixel is an area. They add a level of subjectivity to the data (make more assumptions through the functions than the data can handle). For most applications this is fine, but in scientific applications where detection of the faintest possible galaxies or fainter parts of bright galaxies is our aim, we cannot afford this loss. Because of these reasons Warp will not use such interpolation techniques.

Warp will do interpolation based on "pixel mixing"[28] or "area resampling". This is also what the Hubble Space Telescope pipeline calls "Drizzling"[29]. This technique requires no functions, it is thus non-parametric. It is also the closest we can get (make least assumptions) to what actually happens on the detector pixels. The basic idea is that you reverse-transform each output pixel to find which pixels of the input image it covers and what fraction of the area of the input pixels are covered. To find the output pixel value, you simply sum the value of each input pixel weighted by the overlap fraction (between 0 to 1) of the output pixel and that input pixel. Through this process, pixels are treated as an area not as a point (which is how detectors create the image), also the brightness (see Section 8.1.3 [Flux Brightness and magnitude], page 201) of an object will be left completely unchanged.

If there are very high spatial-frequency signals in the image (for example fringes) which vary on a scale smaller than your output image pixel size, pixel mixing can cause ailiasing[30]. So if the input image has fringes, they have to be calculated and removed separately (which would naturally be done in any astronomical application). Because of the PSF no astronomical target has a sharp change in the signal so this issue is less important for astronomical applications, see Section 8.1.1.2 [Point Spread Function], page 196.

6.4.4 Invoking Warp

Warp an input dataset into a new grid. Any homographic warp (for example scaling, rotation, translation, projection) is acceptable, see Section 6.4.1 [Warping basics], page 139, for the definitions. The general template for invoking Warp is:

```
$ astwarp [OPTIONS...] InputImage
```

One line examples:

```
## Rotate and then scale input image:
$ astwarp --rotate=37.92 --scale=0.8 image.fits

## Scale, then translate the input image:
$ astwarp --scale 8/3 --translate 2.1 image.fits
```

[27] see http://entropymine.com/imageworsener/bicubic/ for a nice introduction.

[28] For a graphic demonstration see http://entropymine.com/imageworsener/pixelmixing/.

[29] http://en.wikipedia.org/wiki/Drizzle_(image_processing)

[30] http://en.wikipedia.org/wiki/Aliasing

```
## Align raw image with celestial coordinates:
$ astwarp --align rawimage.fits --output=aligned.fits

## Directly input a custom warping matrix (using fraction):
$ astwarp --matrix=1/5,0,4/10,0,1/5,4/10,0,0,1 image.fits

## Directly input a custom warping matrix, with final numbers:
$ astwarp --matrix="0.7071,-0.7071,  0.7071,0.7071" image.fits
```

If any processing is to be done, Warp can accept one file as input. As in all Gnuastro programs, when an output is not explicitly set with the `--output` option, the output filename will be set automatically based on the operation, see Section 4.8 [Automatic output], page 77. For the full list of general options to all Gnuastro programs (including Warp), please see Section 4.1.2 [Common options], page 52.

To be the most accurate, the input image will be read as a 64-bit double precision floating point dataset and all internal processing is done in this format (including the raw output type). You can use the common `--type` option to write the output in any type you want, see Section 4.4 [Numeric data types], page 64.

Warps must be specified as command-line options, either as (possibly multiple) modular warpings (for example `--rotate`, or `--scale`), or directly as a single raw matrix (with `--matrix`). If specified together, the latter (direct matrix) will take precedence and all the modular warpings will be ignored. Any number of modular warpings can be specified on the command-line and configuration files. If more than one modular warping is given, all will be merged to create one warping matrix. As described in Section 6.4.2 [Merging multiple warpings], page 142, matrix multiplication is not commutative, so the order of specifying the modular warpings on the command-line, and/or configuration files makes a difference (see Section 4.2.2 [Configuration file precedence], page 60). The full list of modular warpings and the other options particular to Warp are described below.

The values to the warping options (modular warpings as well as `--matrix`), are a sequence of at least one number. Each number in this seqence is separated from the next by a comma (`,`). Each number can also be written as a single fraction (with a forward-slash `/` between the numerator and denomiator). Space and Tab characters are permitted between any two numbers, just don't forget to quote the whole value. Otherwise, the value will not be fully passed onto the option. See the examples above as a demonstration.

Based on the FITS standard, integer values are assigned to the center of a pixel and the coordinate [1.0, 1.0] is the center of the first pixel (bottom left of the image when viewed in SAO ds9). So the coordinate center [0.0, 0.0] is half a pixel away (in each axis) from the bottom left vertice of the first pixel. The resampling that is done in Warp (see Section 6.4.3 [Resampling], page 142) is done on the coordinate axises and thus directly depends on the coordinate center. In some situations this if fine, for example when rotating/aligning a real image, all the edge pixels will be similiarly affected. But in other situations (for example when scaling an over-sampled mock image to its intended resolution, this is not desired: you want the center of the coordinates to be on the corner of the pixel. In such cases, you can use the `--centeroncorner` option which will shift the center by 0.5 before the main warp, then shift it back by -0.5 after the main warp, see below.

`-a`
`--align` Align the image and celestial (WCS) axes given in the input. After it, the vertical image direction (when viewed in SAO ds9) corresponds to the declination and the horizontal axis is the inverse of the Right Ascension (RA). The inverse of the RA is chosen so the image can correspond to what you would actually see on the sky and is common in most survey images.

Align is internally treated just like a rotation (`--rotation`), but uses the input image's WCS to find the rotation angle. Thus, if you have rotated the image before calling `--align`, you might get unexpected results (because the rotation is defined on the original WCS).

`-r FLT`
`--rotate=FLT`
Rotate the input image by the given angle in degrees: θ in Section 6.4.1 [Warping basics], page 139. Note that commonly, the WCS structure of the image is set such that the RA is the inverse of the image horizontal axis which increases towards the right in the FITS standard and as viewed by SAO ds9. So the default center for rotation is on the right of the image. If you want to rotate about other points, you have to translate the warping center first (with `--translate`) then apply your rotation and then return the center back to the original position (with another call to `--translate`, see Section 6.4.2 [Merging multiple warpings], page 142.

`-s FLT[,FLT]`
`--scale=FLT[,FLT]`
Scale the input image by the given factor(s): M and N in Section 6.4.1 [Warping basics], page 139. If only one value is given, then both image axes will be scaled with the given value. When two values are given (separated by a comma), the first will be used to scale the first axis and the second will be used for the second axis. If you only need to scale one axis, use `1` for the axis you don't need to scale. The value(s) can also be written (on the command-line or in configuration files) as a fraction.

`-f FLT[,FLT]`
`--flip=FLT[,FLT]`
Flip the input image around the given axis(s). If only one value is given, then both image axes are flipped. When two values are given (separated by a comma), you can choose which axis to flip over. `--flip` only takes values `0` (for no flip), or `1` (for a flip). Hence, if you want to flip by the second axis only, use `--flip=0,1`.

`-e FLT[,FLT]`
`--shear=FLT[,FLT]`
Shear the input image by the given value(s): A and B in Section 6.4.1 [Warping basics], page 139. If only one value is given, then both image axes will be sheared with the given value. When two values are given (separated by a comma), the first will be used to shear the first axis and the second will be used for the second axis. If you only need to shear along one axis, use `0` for the axis that

must be untouched. The value(s) can also be written (on the command-line or in configuration files) as a fraction.

`-t FLT[,FLT]`
`--translate=FLT[,FLT]`

Translate (move the center of coordinates) the input image by the given value(s): c and f in Section 6.4.1 [Warping basics], page 139. If only one value is given, then both image axes will be translated by the given value. When two values are given (separated by a comma), the first will be used to translate the first axis and the second will be used for the second axis. If you only need to translate along one axis, use `0` for the axis that must be untouched. The value(s) can also be written (on the command-line or in configuration files) as a fraction.

`-p FLT[,FLT]`
`--project=FLT[,FLT]`

Apply a projection to the input image by the given values(s): g and h in Section 6.4.1 [Warping basics], page 139. If only one value is given, then projection will apply to both axises with the given value. When two values are given (separated by a comma), the first will be used to project the first axis and the second will be used for the second axis. If you only need to project along one axis, use `0` for the axis that must be untouched. The value(s) can also be written (on the command-line or in configuration files) as a fraction.

`-m STR`
`--matrix=STR`

The warp/transformation matrix. All the elements in this matrix must be separated by comas(`,`) characters and as described above, you can also use fractions (a forward-slash between two numbers). The transformation matrix can be either a 2 by 2 (4 numbers), or a 3 by 3 (9 numbers) array. In the former case (if a 2 by 2 matrix is given), then it is put into a 3 by 3 matrix (see Section 6.4.1 [Warping basics], page 139).

The determinant of the matrix has to be non-zero and it must not contain any non-number values (for example infinities or NaNs). The elements of the matrix have to be written row by row. So for the general homography matrix of Section 6.4.1 [Warping basics], page 139, it should be called with `--matrix=a,b,c,d,e,f,g,h,1`.

The raw matrix takes precedence over all the modular warping options listed above, so if it is called with any number of modular warps, the latter are ignored.

`-c`
`--centeroncorer`

Put the center of coordinates on the corner of the first (bottom-left when viewed in SAO ds9) pixel. This option is applied after the final warping matrix has been finalized: either through modular warpings or the raw matrix. See the explanation above for coordinates in the FITS standard to better understand this option and when it should be used.

`--hstartwcs=INT`
: Specify the first header keyword number (line) that should be used to read the WCS information, see the full explanation in Section 6.1.4 [Invoking Crop], page 101.

`--hendwcs=INT`
: Specify the last header keyword number (line) that should be used to read the WCS information, see the full explanation in Section 6.1.4 [Invoking Crop], page 101.

`-k`
`--keepwcs`
: Do not correct the WCS information of the input image and save it untouched to the output image. By default the WCS (World Coordinate System) information of the input image is going to be corrected in the output image so the objects in the image are at the same WCS coordinates. But in some cases it might be useful to keep it unchanged (for example to correct alignments).

`-C FLT`
`--coveredfrac=FLT`
: Depending on the warp, the output pixels that cover pixels on the edge of the input image, or blank pixels in the input image, are not going to be fully covered by input data. With this option, you can specify the acceptable covered fraction of such pixels (any value between 0 and 1). If you only want output pixels that are fully covered by the input image area (and are not blank), then you can set `--coveredfrac=1`. Alternatively, a value of `0` will keep output pixels that are even infinitesmially covered by the input(so the sum of the pixels in the input and output images will be the same).

7 Data analysis

Astronomical datasets (images or tables) contain very valuable information, the tools in this section can help in analysing, extracting, and quantifying that information. For example getting general or specific statistics of the dataset (with Section 7.1 [Statistics], page 148), detecting signal within a noisy dataset (with Section 7.2 [NoiseChisel], page 163), or creating a catalog from an input dataset (with Section 7.3 [MakeCatalog], page 175).

7.1 Statistics

The distribution of values in a dataset can provide valuable information about it. For example, in an image, if it is a positively skewed distribution, we can see that there is significant data in the image. If the distribution is roughly symmetric, we can tell that there is no significant data in the image. In a table, when we need to select a sample of objects, it is important to first get a general view of the whole sample.

On the other hand, you might need to know certain statistical parameters of the dataset. For example, if we have run a detection algorithm on an image, and we want to see how accurate it was, one method is to calculate the average of the undetected pixels and see how reasonable it is (if detection is done correctly, the average of undetected pixels should be approximately equal to the background value, see Section 7.1.3 [Sky value], page 150). In a table, you might have calcuated the magnitudes of a certain class of objects and want to get some general characteristics of the distribution immediately on the command-line (very fast!), to possibly change some parameters. The Statistics program is designed for such situations.

7.1.1 Histogram and Cumulative Frequency Plot

Histograms and the cumulative frequency plots are both used to visually study the distribution of a dataset. A histogram shows the number of data points which lie within pre-defined intervals (bins). So on the horizontal axis we have the bin centers and on the vertical, the number of points that are in that bin. You can use it to get a general view of the distribution: which values have been repeated the most? how close/far are the most significant bins? Are there more values in the larger part of the range of the dataset, or in the lower part? Similarly, many very important properties about the dataset can be deduced from a visual inspection of the histogram. In the Statistics program, the histogram can be either output to a table to plot with your favorite plotting program[1], or it can be shown with ASCII characters on the command-line, which is very crude, but good enough for a fast and on-the-go analysis, see the example in Section 7.1.4 [Invoking Statistics], page 154.

The width of the bins is only necessary parameter for a histogram. In the limiting case that the bin-widths tend to zero (while assuming the number of points in the dataset tend to infinity), then the histogram will tend to the probability density function (`https://en.wikipedia.org/wiki/Probability_density_function`) of the distribution. When the absolute number of points in each bin is not relevant to the study (only the shape of the histogram is important), you can *normalize* a histogram so like the probability density function, the sum of all its bins will be one.

---

[1] We recommend PGFPlots (`http://pgfplots.sourceforge.net/`) which generates your plots directly within TEX (the same tool that generates your document).

In the cumulative frequency plot of a distribution, the horizontal axis is the sorted data values and the y axis is the index of each data in the sorted distribution. Unlike a histogram, a cumulative frequency plot does not involve intervals or bins. This makes it less prone to any sort of bias or error that a given bin-width would have on the analysis. When a larger number of the data points have roughly the same value, then the cumulative frequency plot will become steep in that vicinity. This occurs because on the horizontal axis, there is little change while on the vertical axis, the indexes constantly increase. Normalizing a cumulative frequency plot means to divide each index (y axis) by the total number of data points (or the last value).

Unlike the histogram which has a limited number of bins, ideally the cumulative frequency plot should have one point for every data element. Even in small datasets (for example a 200×200 image) this will result in an unreasonably large number of points to plot (40000)! As a result, for practical reasons, it is common to only store its value on a certain number of points (intervals) in the input range rather than the whole dataset, so you should determine the number of bins you want when asking for a cumulative frequency plot. In Gnuastro (and thus the Statistics program), the number reported for each bin is the total number of datapoints until the larger interval value for that bin. You can see an example histogram and cumulative frequency plot of a single dataset under the `--asciihist` and `--asciicfp` options of Section 7.1.4 [Invoking Statistics], page 154.

So as a summary, both the histogram and cumulative frequency plot in Statistics will work with bins. Within each bin/interval, the lower value is considered to be within then bin (it is inclusive), but its larger value is not (it is exclusive). Formally, an interval/bin between a and b is represented by [a, b). When the over-all range of the dataset is specified (with the `--greaterequal`, `--lessthan`, or `--qrange` options), the acceptable values of the dataset are also defined with a similar inclusive-exclusive manner. But when the range is determined from the actual dataset (none of these options is called), the last element in the dataset is included in the last bin's count.

7.1.2 Sigma clipping

Let's assume that you have pure noise (centered on zero) with a clear Gaussian distribution (`https://en.wikipedia.org/wiki/Normal_distribution`), or see Section 8.2.1.1 [Photon counting noise], page 210. Now let's assume you add very bright objects (signal) on the image which have a very sharp boundary. By a sharp boundary, we mean that there is a clear cutoff (from the noise) at the pixels the objects finish. In other words, at their boundaries, the objects do not fade away into the noise. In such a case, when you plot the histogram (see Section 7.1.1 [Histogram and Cumulative Frequency Plot], page 148) of the distribution, the pixels relating to those objects will be clearly separate from pixels that belong to parts of the image that did not have any signal (were just noise). In the cumulative frequency plot, after a steady rise (due to the noise), you would observe a long flat region were for a certain range of data (horizontal axis), there is no increase in the index (vertical axis).

Outliers like the example above can significantly bias the measurement of noise statistics. σ-clipping is defined as a way to avoid the effect of such outliers. In astronomical applications, cosmic rays (when they collide at a near normal incidence angle) are a very good example of such outliers. The tracks they leave behind in the image are perfectly immune to the blurring caused by the atmosphere and the aperture. They are also very

energetic and so their borders are usually clearly separated from the surrounding noise. So σ-clipping is very useful in removing their effect on the data. See Figure 15 in Akhlaghi and Ichikawa, 2015 (`https://arxiv.org/abs/1505.01664`).

σ-clipping is defined as the very simple iteration below. In each iteration, the range of input data might decrease and so when the outliers have the conditions above, the outliers will be removed through this iteration. The exit criteria will be discussed below.

1. Calculate the standard deviation (σ) and median (m) of a distribution.
2. Remove all points that are smaller or larger than $m \pm \alpha\sigma$.
3. Go back to step 1, unless the selected exit criteria is reached.

The reason the median is used as a reference and not the mean is that the mean is too significantly affected by the presence of outliers, while the median is less affected, see Section 7.1.3.3 [Quantifying signal in a tile], page 153. As you can tell from this algorithm, besides the condition above (that the signal have clear high signal to noise boundaries) σ-clipping is only useful when the signal does not cover more than half of the full data set. If they do, then the median will lie over the outliers and σ-clipping might remove the pixels with no signal.

There are commonly two exit criteria to stop the σ-clipping iteration:

- When a certain number of iterations has taken place (second value to the `--sigclip` option is larger than 1).
- When the new measured standard deviation is within a certain tolerance level of the old one (second value to the `--sigclip` option is less than 1). The tolerance level is defined by:

  $$\frac{\sigma_{old} - \sigma_{new}}{\sigma_{new}}$$

  The standard deviation is used because it is heavily influenced by the presence of outliers. Therefore the fact that it stops changing between two iterations is a sign that we have successfully removed outliers. Note that in each clipping, the dispersion in the distribution is either less or equal. So $\sigma_{old} \geq \sigma_{new}$.

> When working on astronomical images, objects like galaxies and stars are blurred by the atmosphere and the telescope aperture, therefore their signal sinks into the noise very gradually. Galaxies in particular do not appear to have a clear high signal to noise cutoff at all. Therefore σ-clipping will not be useful in removing their effect on the data.
>
> To guage if σ-clipping will be useful for your dataset, look at the histogram (see Section 7.1.1 [Histogram and Cumulative Frequency Plot], page 148). The ASCII histogram that is printed on the command-line with `--asciihist` is good enough in most cases.

7.1.3 Sky value

One of the most important aspects of a dataset is its reference value: the value of the dataset where there is no signal. Without knowing, and thus removing the effect of, this value it is impossible to compare the derived results of many high-level analyses over the

dataset with other datasets (in the attempt to associate our results with the "real" world). In astronomy, this reference value is known as the "Sky" value: the value where there is no signal from objects (for example galaxies, stars, planets or comets). Depending on the dataset, the Sky value maybe a fixed value over the whole dataset, or it may vary based on location. For an example of the latter case, see Figure 11 in Akhlaghi and Ichikawa (2015) (`https://arxiv.org/abs/1505.01664`).

Because of the significance of the Sky value in astronomical data analysis, we have devoted this subsection to it for a thorough review. We start with a thorough discussion on its definition (Section 7.1.3.1 [Sky value definition], page 151). In the astronomical literature, researchers use a variety of methods to estimate the Sky value, so in Section 7.1.3.2 [Sky value misconceptions], page 152) we review those and discuss their biases. From the definition of the Sky value, the most accurate way to estimate the Sky value is to run a detection algorithm (for example Section 7.2 [NoiseChisel], page 163) over the dataset and use the un-detected pixels. However, there is also a more crude method that maybe useful when good direct detection is not initially possible (for example due to too many cosmic rays in a shallow image). A more crude (but simpler method) that is usable in such situations is discussed in Section 7.1.3.3 [Quantifying signal in a tile], page 153.

7.1.3.1 Sky value definition

This analysis is taken from Akhlaghi and Ichikawa (2015) (`https://arxiv.org/abs/1505.01664`). Let's assume that all instrument defects – bias, dark and flat – have been corrected and the brightness (see Section 8.1.3 [Flux Brightness and magnitude], page 201) of a detected object, O, is desired. The sources of flux on pixel i[2] of the image can be written as follows:

- Contribution from the target object, (O_i).
- Contribution from other detected objects, (D_i).
- Undetected objects or the fainter undetected regions of bright objects, (U_i).
- A cosmic ray, (C_i).
- The background flux, which is defined to be the count if none of the others exists on that pixel, (B_i).

The total flux in this pixel (T_i) can thus be written as:

$$T_i = B_i + D_i + U_i + C_i + O_i.$$

By definition, D_i is detected and it can be assumed that it is correctly estimated (deblended) and subtracted, thus $D_i = 0$. There are also methods to detect and remove cosmic rays, for example the method described in van Dokkum (2001)[3], or by comparing multiple exposures. This allows us to set $C_i = 0$. Note that in practice, D_i and U_i are correlated, because they both directly depend on the detection algorithm and its input parameters. Also note that

[2] For this analysis the dimension of the data (image) is irrelevant. So if the data is an image (2D) with width of w pixels, then a pixel located on column x and row y (where all counting starts from zero and (0, 0) is located on the bottom left corner of the image), would have an index: $i = x + y \times w$.

[3] van Dokkum, P. G. (2001). Publications of the Astronomical Society of the Pacific. 113, 1420.

no detection or cosmic ray removal algorithm is perfect. With these limitations in mind, the observed Sky value for this pixel (S_i) can be defined as

$$S_i = B_i + U_i.$$

Therefore, as the detection process (algorithm and input parameters) becomes more accurate, or $U_i \to 0$, the sky value will tend to the background value or $S_i \to B_i$. Therefore, while B_i is an inherent property of the data (pixel in an image), S_i depends on the detection process. Over a group of pixels, for example in an image or part of an image, this equation translates to the average of undetected pixels. With this definition of sky, the object flux in the data can be calculated with

$$T_i = S_i + O_i \quad \to \quad O_i = T_i - S_i.$$

Hence, the more accurately S_i is measured, the more accurately the brightness (sum of pixel values) of the target object can be measured (photometry). Any under-(over-)estimation in the sky will directly translate to an over-(under-)estimation of the measured object's brightness. In the fainter outskirts of an object a very small fraction of the photo-electrons in the pixels actually belong to objects (see Figure 1b in Akhlaghi and Ichikawa (2015) (`https://arxiv.org/abs/1505.01664`)). Therefore even a small over estimation of the sky value will result in the loss of a very large portion of most galaxies. Besides the lost area/brightness, this will also cause an over-estimation of the Sky value and thus even more under-estimation of the object's brightness. It is thus very important to detect the diffuse flux of a target, even if they are not your primary target.

The **Sky value** is only correctly found when all the detected objects (D_i and C_i) have been removed from the data.

7.1.3.2 Sky value misconceptions

As defined in Section 7.1.3 [Sky value], page 150, the sky value is only accurately defined when the detection algorithm is not significantly reliant on the sky value. In particular its detection threshold. However, most signal-based detection tools[4] use the sky value as a reference to define the detection threshold. So these old techniques had to rely on approximations based on other assumptions about the data. A review of those other techniques can be seen in Appendix A of Akhlaghi and Ichikawa (2015)[5]. Since they were extensively used in astronomical data analysis for several decades, such approximations have given rise to a lot of misconceptions, ambiguities and disagreements about the sky value and how to measure it. As a summary, the major methods used until now were an approximation of the mode of the image pixel distribution and σ-clipping.

- To find the mode of a distribution those methods would either have to assume (or find) a certain probability density function (PDF) or use the histogram. But astronomical

[4] According to Akhlaghi and Ichikawa (2015), signal-based detection is a detection process that relies heavily on assumptions about the to-be-detected objects. This method was the most heavily used technique prior to the introduction of NoiseChisel in that paper.

[5] Akhlaghi M., Ichikawa. T. (2015). Astrophysical Journal Supplement Series.

datasets can have any distribution, making it almost impossible to define a generic function. Also, histogram-based results are very inaccurate (there is a large dispersion) and it depends on the histogram bin-widths.

- Another approach was to iteratively clip the brightest pixels in the image (which is known as σ-clipping, since the reference was found from the image mean and its standard deviation or σ). See Section 7.1.2 [Sigma clipping], page 149, for a complete explanation. The problem with σ-clipping was that real astronomical objects have diffuse and faint wings that penetrate deeply into the noise. So only removing their brightest parts is completely useless in removing the systematic bias an object's fainter parts cause in the sky value.

As discussed in Section 7.1.3 [Sky value], page 150, the sky value can only be correctly defined as the average of undetected pixels. Therefore all such approaches that try to approximate the sky value prior to detection are ultimately poor approximations.

7.1.3.3 Quantifying signal in a tile

Put simply, noise can characterized with a certain spread about a characteristic value. In the Gaussian distribution (most commonly used to model noise) the spread is defined by the standard deviation about the characteristic mean. Before continuing let's clarify some definitions first: *Data* is defined as the combination of signal and noise (so a noisy image is one *data*-set). *Signal* is defined as the mean of the noise on each element (after sky subtraction, see Section 7.1.3.1 [Sky value definition], page 151).

Let's assume that the *background* (see Section 7.1.3.1 [Sky value definition], page 151) is subtracted and is zero. When a data set doesn't have any signal (only noise), the mean, median and mode of the distribution are equal within statistical errors and approximately equal to the background value. Signal always has a positive value and will never become negative, see Figure 1 in Akhlaghi and Ichikawa (2015) (`https://arxiv.org/abs/1505.01664`). Therefore, as more signal is added to the raw noise, the mean, median and mode of the dataset (which has both signal and noise) shift to the positive. The mean's shift is the largest. The median shifts less, since it is defined based on an ordered distribution and so is not affected by a small number of outliers. The distribution's mode shifts the least to the positive.

Inverting the argument above gives us a robust method to quantify the significance of signal in a dataset. Namely, when the mode and median of a distribution are approximately equal, we can argue that there is no significant signal. To allow for gradients (which are commonly present in ground-based images), we can consider the image to be made of a grid of tiles (see Section 4.6 [Tessellation], page 72[6]). Hence, from the difference of the mode and median on each tile, we can 'detect' the significance of signal in it. The median of a distribution is defined to be the value of the distribution's middle point after sorting (or 0.5 quantile). Thus, to estimate the presence of signal, we'll compare with the quantile of the mode with 0.5, if the difference is larger than the value given to the `--modmedqdiff` option, this tile will be ignored. You can read this option as "mode-median-quantile-diff".

This method to use the input's skewness is possible because of a new algorithm to find the mode of a distribution that was defined in Appendix C of Akhlaghi and Ichikawa (2015). However, the raw dataset's distribution is noisy (noise also affects the sorting), so using the

[6] The options to customize the tessellation are discussed in Section 4.1.2.2 [Processing options], page 54.

argument above on the raw input will give a noisy result. To decrease the noise/error in estimating the mode, we will use convolution (see Section 6.3.1.1 [Convolution process], page 118). Convolution decreases the range of the dataset and enhances its skewness, See Section 3.1.1 and Figure 4 in Akhlaghi and Ichikawa (2015). This enhanced skewness can be interpreted as an increase in the Signal to noise ratio of the objects buried in the noise. Therefore, to obtain an even better measure of the presence of signal in a mesh, the image can be convolved with a given kernel first.

Note that through the difference of the mode and median we have actually 'detected' data in the distribution. However this "detection" was only based on the total distribution of the data in each tile (a much lower resolution). This is the main limitation of this technique. The best approach is thus to do detection over the dataset, mask all the detected pixels and use the undetected regions to estimate the sky and its standard deviation.

The mean value of the tiles that have an approximately equal mode and median will be the Sky value. However there is one final hudle: astronomical datasets are commonly plagued with Cosmic rays. Images of Cosmic rays aren't smoothed by the atmosphere or telescope aperture, so they have sharp boundaries. Also, since they don't occupy too many pixels, they don't affect the mode and median calculation. But their very high values can greatly bias the calculation of the mean (recall how the mean shifts the fastest in the presence of outliers), see Figure 15 in Akhlaghi and Ichikawa (2015) for one example.

The effect of outliers like cosmic rays on the mean and standard deviation can be removed through σ-clipping, see Section 7.1.2 [Sigma clipping], page 149, for a complete explanation. Therefore, after asserting that the mode and median are approximately equal in a tile (see Section 4.6 [Tessellation], page 72), the final Sky value and its standard deviation are determined after σ-clipping with the `--sigmaclip` option.

7.1.4 Invoking Statistics

Statistics will print statistical measures of an input dataset (table column or image). The executable name is `aststatistics` with the following general template

```
$ aststatistics [OPTION ...] InputImage.fits
```

One line examples:

```
## Print some general statistics of input image:
$ aststatistics image.fits

## Print some general statistics of column named MAG_F160W:
$ aststatistics catalog.fits -h1 --column=MAG_F160W

## Make the histogram of the column named MAG_F160W:
$ aststatistics table.fits -cMAG_F160W --histogram

## Find the Sky value on image with a given kernel:
$ aststatistics image.fits --sky --kernel=kernel.fits

## Print Sigma-clipped results of records with a MAG_F160W
## column value between 26 and 27:
$ aststatistics cat.fits -cMAG_F160W -g26 -l27 --sigmaclip=3,0.2
```

```
## Print the median value of all records in column MAG_F160W that
## have a value larger than 3 in column PHOTO_Z:
$ aststatistics tab.txt -rPHOTO_Z -g3 -cMAG_F160W --median
```

An input image or table is necessary when processing is to be done. If any output file is to be created, the value to the `--output` option, is used as the base name for the generated files. Without `--output`, the input name will be used to generate an output name, see Section 4.8 [Automatic output], page 77. The options described below are particular to Statistics, but for general operations, it shares a large collection of options with the other Gnuastro programs, see Section 4.1.2 [Common options], page 52, for the full list. Options can also be given in configuration files, for more, please see Section 4.2 [Configuration files], page 59.

The input dataset may have blank values (see Section 6.1.3 [Blank pixels], page 100), in this case, all blank pixels are ignored during the calculation. Initially, the full dataset will be read, but it is possible to select a specific range of data elements to use in the analysis of each run. You can either directly specify a minimum and maximum value for the range of data elements to use (with `--greaterequal` or `--lessthan`), or specify the range using quantiles (with `--qrange`). If a range is specified, all pixels outside of it are ignored before any processing.

The following set of options are for specifying the input/outputs of Statistics. There are many other input/output options that are common to all Gnuastro programs including Statistics, see Section 4.1.2.1 [Input/Output options], page 52, for those.

`-c STR/INT`
`--column=STR/INT`
: The input column selector when the input file is a table. See Section 4.5.3 [Selecting table columns], page 71, for a full description of how to use this option. For more on how tables are read in Gnuastro, please see Section 4.5 [Tables], page 66.

`-r STR/INT`
`--refcol=STR/INT`
: The reference column selector when the input file is a table. When a reference column is given, the range options below will be applied to this column and only elements in the input column that have a reference value in the correct range will be used. In practice this option allows you to select a subset of the input column based on values in another (the reference) column. All the statistical calculations will be done on the selected input column, not the reference column.

`-g FLT`
`--greaterequal=FLT`
: Limit the range of inputs into those with values greater and equal to what is given to this option. None of the values below this value will be used in any of the processing steps below.

`-l FLT`
`--lessthan=FLT`

Limit the range of inputs into those with values less-than what is given to this option. None of the values greater or equal to this value will be used in any of the processing steps below.

`-Q FLT[,FLT]`
`--qrange=FLT[,FLT]`

Specify the range of usable inputs using the quantile. This option can take one or two quantiles to specify the range. When only one number is input (let's call it Q), the range will be those values in the quantile range Q to $1 - Q$. So when only one value is given, it must be less than 0.5. When two values are given, the first is used as the lower quantile range and the second is used as the larger quantile range.

The quantile of a given element in a dataset is defined by the fraction of its index to the total number of values in the sorted input array. So the smallest and largest values in the dataset have a quantile of 0.0 and 1.0. The quantile is a very useful non-parametric (making no assumptions about the input) relative measure to specify a range. It can best be understood in terms of the cumulative frequency plot, see Section 7.1.1 [Histogram and Cumulative Frequency Plot], page 148. The quantile of each horizontal axis value in the cumulative frequency plot is the vertical axis value associate with it.

When no operation is requested, Statistics will print some general basic properties of the input dataset on the command-line like the example below (ran on one of the output images of `make check`[7]). This default behavior is designed to help give you a general feeling of how the data are distributed and help in narrowing down your analysis.

```
$ aststatistics convolve_spatial_scaled_noised.fits     \
                --greaterequal=9500 --lessthan=11000
Statistics (GNU Astronomy Utilities) X.X
-------
Input: convolve_spatial_scaled_noised.fits (hdu: 0)
Range: from (inclusive) 9500, upto (exclusive) 11000.
Unit: Brightness
-------
  Number of elements:                         9074
  Minimum:                                    9622.35
  Maximum:                                    10999.7
  Mode:                                       10055.45996
  Mode quantile:                              0.4001983908
  Median:                                     10093.7
  Mean:                                       10143.98257
  Standard deviation:                         221.80834
-------
Histogram:
```

[7] You can try it by running the command in the `tests` directory, open the image with a FITS viewer and have a look at it to get a sence of how these statistics relate to the input image/dataset.

```
|                   **
|                 ******
|                 *******
|                ********* 
|              *************
|              **************
|            ******************
|            ********************
|          *************************** *
|        ***************************************** ***
|*  **************************************************************
|-----------------------------------------------------------------
```

Gnuastro's Statistics is a very general purpose program, so to be able to easily understand this diversity in its operations (and how to possibly run them together), we'll devided the operations into two types: those that don't respect the position of the elements and those that do (by tessellating the input on a tile grid, see Section 4.6 [Tessellation], page 72). The former treat the whole dataset as one and can re-arrange all the elements (for example sort them), but the former do their processing on each tile independently. First, we'll review the operations that work on the whole dataset.

The group of options below can used to get a single value measurement of the whole dataset. They will print only the requested value as one field in a line/row, like the `--mean`, `--median` options. These options can be called any number of times and in any order. The outputs of all such options will be printed on one line following each other (with a space character between them). This feature makes these options very useful in scripts, or to redirect into programs like GNU AWK for higher-level processing. These are some of the most basic measures, Gnuastro is still under heavy development and this list will grow. If you want another statistical parameter, please contact us and we will do out best to add it to this list, see Section 1.8 [Suggest new feature], page 11.

`-n`
`--number` Print the number of all used (non-blank and in range) elements.

`--minimum`
Print the minimum value of all used elements.

`--maximum`
Print the maximum value of all used elements.

`--sum` Print the sum of all used elements.

`-m`
`--mean` Print the mean (average) of all used elements.

`-t`
`--std` Print the standard deviation of all used elements.

`-E`
`--median` Print the median of all used elements.

`-u FLT[,FLT[,...]]`
`--quantile=FLT[,FLT[,...]]`

Print the values at the given quantiles of the input dataset. Any number of quantiles may be given and one number will be printed for each. Values can either be written as a single number or as fractions, but must be between zero and one (inclusive). Hence, in effect `--quantile=0.25 --quantile=0.75` is equivalent to `--quantile=0.25,3/4`, or `-u1/4,3/4`.

The returned value is one of the elements from the dataset. Taking q to be your desired quantile, and N to be the total number of used (non-blank and within the given range) elements, the returned value is at the following position in the sorted array: $round(q \times N)$.

`--quantfunc=FLT[,FLT[,...]]`

Print the quantiles of the given values in the dataset. This option is the inverse of the `--quantile` and operates similarly except that the acceptable values are within the range of the dataset, not between 0 and 1. Formally it is known as the "Quantile function".

Since the dataset is not continuous this function will find the nearest element of the dataset and use its position to estimate the quantile function.

`-O`
`--mode`

Print the mode of all used elements. The mode is found through the mirror distribution which is fully described in Appendix C of Akhlaghi and Ichikawa 2015 (`https://arxiv.org/abs/1505.01664`). See that section for a full description.

This mode calculation algorithm is non-parametric, so when the dataset is not large enough (larger than about 1000 elements usually), or doesn't have a clear mode it can fail. In such cases, this option will return a value of `nan` (for the floating point NaN value).

As described in that paper, the easiest way to assess the quality of this mode calculation method is to use it's symmetricity (see `--modesym` below). A better way would be to use the `--mirror` option to generate the histogram and cumulative frequency tables for any given mirror value (the mode in this case) as a table. If you generate plots like those shown in Figure 21 of that paper, then your mode is accurate.

`--modequant`

Print the quantile of the mode. You can get the actual mode value from the `--mode` described above. In many cases, the absolute value of the mode is irrelevant, but its position within the distribution is important. In such cases, this option will become handy.

`--modesym`

Print the symmetricity of the calculated mode. See the description of `--mode` for more. This mode algorithm finds the mode based on how symmetric it is, so if the symmetricity returned by this option is too low, the mode is not too accurate. See Appendix C of Akhlaghi and Ichikawa 2015 (`https://arxiv.org/abs/1505.01664`) for a full description. In practice, symmetricity values larger than 0.2 are mostly good.

`--modesymvalue`
: Print the value in the distribution where the mirror and input distributions are no longer symmetric, see `--mode` and Appendix C of Akhlaghi and Ichikawa 2015 (`https://arxiv.org/abs/1505.01664`) for more.

The list of options below are for those statistical operations that output more than one value. So while they can be called together in one run, their outputs will be distinct (each one's output will usually be printed in more than one line).

`-A`
`--asciihist`
: Print an ASCII histogram of the usable values within the input dataset along with some basic information like the example below (from the UVUDF catalog[8]). The width and height of the histogram (in units of character widths and heights on your command-line terminal) can be set with the `--numasciibins` (for the width) and `--asciiheight` options.

  For a full description of the histogram, please see Section 7.1.1 [Histogram and Cumulative Frequency Plot], page 148. An ASCII plot is certainly very crude and cannot be used in any publication, but it is very useful for getting a general feeling of the input dataset very fast and easily on the command-line without having to take your hands off the keyboard (which is a major distraction!). If you want to try it out, you can write it all in one line and ignore the \ and extra spaces.

```
$ aststatistics uvudf_rafelski_2015.fits.gz --hdu=1         \
                --column=MAG_F160W --lessthan=40            \
                --asciihist --numasciibins=55
ASCII Histogram:
Number: 8593
Y: (linear: 0 to 660)
X: (linear: 17.7735 -- 31.4679, in 55 bins)
 |                                         ****
 |                                        *****
 |                                       ******
 |                                      ********
 |                                      *********
 |                                    ***********
 |                                  **************
 |                                *****************
 |                           ***********************
 |                    ********************************
 |*** ***************************************************
 |-------------------------------------------------------
```

`--asciicfp`
: Print the cumulative frequency plot of the usable elements in the input dataset. Please see descriptions under `--asciihist` for more, the example below is from

[8] `https://asd.gsfc.nasa.gov/UVUDF/uvudf_rafelski_2015.fits.gz`

the same input table as that example. To better understand the cumulative frequency plot, please see Section 7.1.1 [Histogram and Cumulative Frequency Plot], page 148.

```
$ aststatistics uvudf_rafelski_2015.fits.gz --hdu=1         \
                --column=MAG_F160W --lessthan=40            \
                --asciicfp --numasciibins=55
ASCII Cumulative frequency plot:
Y: (linear: 0 to 8593)
X: (linear: 17.7735 -- 31.4679, in 55 bins)
 |                                                *******
 |                                             **********
 |                                            ***********
 |                                          *************
 |                                         **************
 |                                        ***************
 |                                      *****************
 |                                    *******************
 |                                ***********************
 |                         ******************************
 |*******************************************************
 |-------------------------------------------------------
```

-H

--histogram

Save the histogram of the usable values in the input dataset into a table. The first column is the value at the center of the bin and the second is the number of points in that bin. If the --cumulative option is also called with this option in a run, then the table will have three columns (the third is the cumulative frequency plot). Through the --numbins and --lowerbin you can modify the first column values and with --normalize and --maxbinone you can modify the second columns. See below for the description of each.

By default (when no --output is specified) a plain text table will be created, see Section 4.5.2 [Gnuastro text table format], page 69. If a FITS name is specified, you can use the common option --tableformat to have it as a FITS ASCII or FITS binary format, see Section 4.1.2 [Common options], page 52. This table can then be fed into your favorite plotting tool and get a much more clean and nice histogram than what the raw command-line can offer you (with the --asciihist option).

-C

--cumulative

Save the cumulative frequency plot of the usable values in the input dataset into a table, similar to --histogram.

-s

--sigmaclip

Do σ-clipping on the usable pixels of the input dataset. See Section 7.1.2 [Sigma clipping], page 149, for a full description on σ-clipping and also to

better understand this option. The σ-clipping parameters can be set through the `--sclipparams` option (see below).

`--mirror=FLT`
: Make a histogram and cumulative frequency plot of the mirror distribution for the given dataset when the mirror is located at the value to this option. The mirror distribution is fully described in Appendix C of Akhlaghi and Ichikawa 2015 (`https://arxiv.org/abs/1505.01664`) and currently it is only used to calculate the mode (see `--mode`).

  Just note that the mirror distribtution is a discrete distribution like the input, so while you may give any number as the value to this option, the actual mirror value is the closest number in the input dataset to this value. If the two numbers are different, Statistics will warn you of the actual mirror value used.

  This option will make a table as output. Depending on your selected name for the output, it will be either a FITS table or a plain text table (which is the default). It contains three columns: the first is the center of the bins, the second is the histogram (with the largest value set to 1) and the third is the normalized cumulative frequency plot of the mirror distribution. The bins will be positioned such that the mode is on the starting interval of one of the bins to make it symmetric around the mirror. With this output file and the input histogram (that you can generate in another run of Statistics, using the `--onebinvalue`), it is possible to make plots like Figure 21 of Akhlaghi and Ichikawa 2015 (`https://arxiv.org/abs/1505.01664`).

The list of options below allow customization of the histogram and cumulative frequency plots (for the `--histogram`, `--cumulative`, `--asciihist`, and `--asciicfp` options).

`--numbins`
: The number of bins (rows) to use in the histogram and the cumulative frequency plot tables (outputs of `--histogram` and `--cumulative`).

`--numasciibins`
: The number of bins (characters) to use in the ASCII plots when printing the histogram and the cumulative frequency plot (outputs of `--asciihist` and `--asciicfp`).

`--asciiheight`
: The number of lines to use when printing the ASCII histogram and cumulative frequency plot on the command-line (outputs of `--asciihist` and `--asciicfp`).

`-n`
`--normalize`
: Normalize the histogram or cumulative frequency plot tables (outputs of `--histogram` and `--cumulative`). For a histogram, the sum of all bins will become one and for a cumulative frequency plot the last bin value will be one.

`--maxbinone`
: Divide all the histogram values by the maximum bin value so it becomes one and the rest are similarly scaled. In some situations (for example if you want to plot the histogram and cumulative frequency plot in one plot) this can be very useful.

`--onebinstart=FLT`
: Make sure that one bin starts with the value to this option. In practice, this will shift the bins used to find the histogram and cumulative frequency plot such that one bin's lower interval becomes this value. For example when the histogram range includes negative and positive values and zero has a special significance in your analysis, then zero will be somewhere in one bin and will mix counts of positive and negative. By setting `--onebinstart=0`, you can make sure that the viewers of the histogram will not be confused without doing the math of setting a range and number of bins.

  Note that by default, the first row of the histogram and cumulative frequency plot show the central values of each bin. So in the example above you will not see the 0.000 in the first column, you will see two symmetric values. If the value is not within the usable input range, this option will be ignored.

All the options described until now were from the first class of operations discusssed above: those that treat the whole dataset as one. However. It often happens that the relative position of the dataset elements over the dataset is significant. For example you don't want one median value for the whole input image, you want to know how the median changes over the image. For such operations, the input has to be tessellated (see Section 4.6 [Tessellation], page 72). Thus this class of options can't currently be called along with the options above in one run of Statistics.

`-t`
`--ontile`
: Do the respective single-valued calculation over one tile of the input dataset, not the whole dataset. This option must be called with atleast one of the single valued options discussed above (for example `--mean` or `--quantile`). The output will be a file in the same format as the input. If the `--oneelempertile` option is called, then one element/pixel will be used for each tile (see Section 4.1.2.2 [Processing options], page 54). Otherwise, the output will have the same size as the input, but each element will have the value corresponding to that tile's value. If multiple single valued operations are called, then for each operation there will be one extension in the output FITS file.

`-y`
`--sky`
: Estimate the Sky value on each tile as fully described in Section 7.1.3.3 [Quantifying signal in a tile], page 153. As described in that section, several options are necessary to configure the Sky estimation which are listed below. The output file will have two extensions: the first is the Sky value and the second is the Sky standard deviation on each tile. Similar to `--ontile`, if the `--oneelempertile` option is called, then one element/pixel will be used for each tile (see Section 4.1.2.2 [Processing options], page 54).

The parameters for estimating the sky value can be set with the following options, except for the `--sclipparams` option (which is also used by the `--sigmaclip`), the rest are only used for the Sky value estimation.

`-k=STR`
`--kernel=STR`
: File name of kernel to help in estimating the significance of signal in a tile, see Section 7.1.3.3 [Quantifying signal in a tile], page 153.

`--khdu=STR`
: Kernel HDU to help in estimating the significance of signal in a tile, see Section 7.1.3.3 [Quantifying signal in a tile], page 153.

`--mirrordist=FLT`
: Maximum distance (as a multiple of error) to estimate the difference between the input and mirror distributions in finding the mode, see Appendix C of Akhlaghi and Ichikawa 2015 (`https://arxiv.org/abs/1505.01664`), also see Section 7.1.3.3 [Quantifying signal in a tile], page 153.

`--modmedqdiff=FLT`
: The maximum acceptable distance between the mode and median, see Section 7.1.3.3 [Quantifying signal in a tile], page 153.

`--sclipparams=FLT,FLT`
: The σ-clipping parameters, see Section 7.1.2 [Sigma clipping], page 149. This option takes two values which are separated by a comma (`,`). Each value can either be written as a single number or as a fraction of two numbers (for example `3,1/10`). The first value to this option is the multiple of σ that will be clipped (α in that section). The second value is the exit criteria. If it is less than 1, then it is interpretted as tolerance and if it is larger than one it is a specific number. Hence, in the latter case the value must be an integer.

`--smoothwidth=INT`
: Width of a flat kernel to convolve the interpolated tile values. Tile interpolation is done using the median of the `--interpnumngb` neighbors of each tile (see Section 4.1.2.2 [Processing options], page 54). If this option is given a value of zero or one, no smoothing will be done. Without smoothing, strong boundaries will probably be created between the values estimated for each tile. It is thus good to smooth the interpolated image so strong discontinuities do not show up in the final Sky values. The smoothing is done through convolution (see Section 6.3.1.1 [Convolution process], page 118) with a flat kernel, so the value to this option must be an odd number.

`--checksky`
: Create a multi-extension FITS file showing the steps that were used to estimate the Sky value over the input, see Section 7.1.3.3 [Quantifying signal in a tile], page 153. The file will have two extensions for each step (one for the Sky and one for the Sky standard deviation).

7.2 NoiseChisel

Once instrumental signatures are removed from the raw data in the initial reduction process (see Chapter 6 [Data manipulation], page 97). We are ready to derive scientific results out of them. But we can't do anything special with a raw dataset, for example an image is just an array of values. Every pixel just has one value and its position within the image. Therefore, the first step of your high-level analysis will be to classify/label the dataset elements/pixels into two classes: signal and noise. This process is formally known as *detection*. Afterwards, you want to separate the detections into multiple components (for example when two detected regions aren't touching, they should be treated independently

as two distant galaxies for example). This higher level classification of the detections is known as *segmentation*. NoiseChisel is Gnuastro's program for detection and segmentation.

NoiseChisel works based on a new noise-based approach to signal detection and was introduced to the astronomical community in Akhlaghi and Ichikawa [2015] (`https://arxiv.org/abs/1505.01664`). NoiseChisel's primary output is an array (image) with the same size as the input but containing labels: those pixels with a label of 0 are noise/sky while those pixels with labels larger than 0 are detections (separate segments will be given positive integers, starting from 1). For more on NoiseChisel's particular output format and its benefits (especially in conjunction with Section 7.3 [MakeCatalog], page 175), please see Akhlaghi [2016] (`https://arxiv.org/abs/1611.06387`). The published paper cannot under go any updates, but the NoiseChisel software has evolved, you can see the major changes in Section 7.2.1 [NoiseChisel changes after publication], page 164.

Data is inherently mixed with noise: only mock/simulated datasets are free of noise. So this process of separating signal from noise is not trivial. In particular, most scientifically interesting astronomical targets are faint, can have a large variety of morphologies along with a large distribution in brightness and size which are all drowned in a ocean of noise. So detection is a uniquely vital aspect of any scientific work and even more so for astronomical research. This is such a fundamental step that designing of NoiseChisel was the primary motivation behind creating Gnuastro: the first generation of Gnuastro's programs were all first part of what later became NoiseChisel, afterwards they spinned-off into separate programs.

The name of NoiseChisel is derived from the first thing it does after thresholding the dataset: to erode it. In mathematical morphology, erosion on pixels can be pictured as carving off boundary pixels. Hence, what NoiseChisel does is similar to what a wood chisel or stone chisel do. It is just not a hardware, but a software. Infact looking at it as a chisel and your dataset as a solid cube of rock will greatly help in best using it: with NoiseChisel you literally carve the galaxies/stars/comets out of the noise. Try running it with the `--checkdetection` option to see each step of the carving process on your input dataset. You can then change a specific option to carve out your signal out of the noise more successfully.

7.2.1 NoiseChisel changes after publication

Before using NoiseChisel it is strongly recommended to read Akhlaghi and Ichikawa [2015] (`https://arxiv.org/abs/1505.01664`) to gain a good understanding of what it does and how each parameter influences the output. Thanks to that paper, there is no more need to continue this introduction any further and we can just dive into the details of running NoiseChisel in Section 7.2.2 [Invoking NoiseChisel], page 165. However, the paper cannot undergo any further updates, but NoiseChisel will evolve: better algorithms or steps will be found, thus options will be added or removed. So this book is the final and definitive guide. To make the transition form the paper to this book easier (and encourage reading the paper), below you can see the major changes since that paper was published.

- `--noerodequant`: to specify a quantile threshold where erosion will not apply. This is useful to detect sharper point-like sources that will be missed due to too much erosion. To completely ignore this features give this option a value of 1 (only the largest valued pixel in the input will not be eroded).

- `--cleandilated`: After dilation, if the signal-to-noise ratio of a detection is less than the derived pseudo-detection S/N limit, that detection will be discarded. In an ideal/clean noise, a true detection's S/N should be larger than its constituent pseudo-detections because its area is larger and it also covers more signal. However, on a false detections (especially at lower `--detquant` values), the increase in size can cause a decrease in S/N below that threshold.

  This will improve purity and not change completeness (a true detection will not be discarded). Because a true detection has flux in its vicinity and dilation will catch more of that flux and increase the S/N. So on a true detection, the final S/N cannot be less than pseudo-detections.

  However, in many real images bad processing creates artifacts that cannot be accurately removed by the Sky subtraction. In such cases, this option will decrease the completeness (will artificially discard true detections). So this feature is not default and should to be explicitly called when you know the noise is clean.

For a more detailed list of updates in each release, please follow the `NEWS` file. The `NEWS` file is in the released Gnuastro tarball (see Section 3.2.1 [Release tarball], page 31). You can also see it online at `http://git.savannah.gnu.org/cgit/gnuastro.git/plain/NEWS`.

7.2.2 Invoking NoiseChisel

NoiseChisel will detect and segment signal in noise producing a multi-extension labeled image, ready for input into Section 7.3 [MakeCatalog], page 175, to generate a catalog or other processing. The executable name is `astnoisechisel` with the following general template

```
$ astnoisechisel [OPTION ...] InputImage.fits
```

One line examples:

```
## Detect signal in input.fits:
$ astnoisechisel input.fits

## Detect signal assuming input has 4 channels along first dimension
## and 1 along the second. Also set the regular tile size to 100 along
## both dimensions:
$ astnoisechisel --numchannels=4,1 --tilesize=100,100 input.fits
```

If NoiseChisel is to do processing (for example you don't want to get help, or see the values to each input parameter), an input image should be provided with the recognized extensions (see Section 4.1.1.1 [Arguments], page 49). NoiseChisel shares a large set of common operations with other Gnuastro programs, mainly regarding input/output, general processing steps, and general operating modes. To help in a unified experience between all of Gnuastro's programs, these operations have the same command-line options, see Section 4.1.2 [Common options], page 52, for a full list. Since the common options are thoroughly discussed there, they are no longer reviewed here. You can see all the options with a short description on the command-line with the `--help` option, see Section 4.7 [Getting help], page 74.

NoiseChisel's input image may contain blank elements (see Section 6.1.3 [Blank pixels], page 100). Blank elements will be ignored in all steps of NoiseChisel. Hence if your dataset has bad pixels which should be masked with a mask image, please use Gnuastro's Section 6.2

[Arithmetic], page 108, program (in particular its `where` operator) to convert those pixels to blank pixels before running NoiseChisel. Gnuastro's Arithmetic program has bitwise operators helping you select specific kinds of bad-pixels when necessary.

A convolution kernel can also be optionally given. If a value (file name) is given to `--kernel` on the command-line or in a configuration file (see Section 4.2 [Configuration files], page 59), then that file will be used to convolve the image prior to thresholding. Otherwise a default kernel will be used. The default kernel is a 2D Gaussian with a FWHM of 2 pixels truncated at 5 times the FWHM. This choice of the default kernel is discussed in Section 3.1.1 of Akhlaghi and Ichikawa [2015] (`https://arxiv.org/abs/1505.01664`). See Section 6.3.4 [Convolution kernel], page 135, for kernel related options.

NoiseChisel defines two tessellations over the input (see Section 4.6 [Tessellation], page 72). This enables it to deal with possible gradients in the input dataset and also significantly improve speed by processing each tile on different threads. The tessellation related options are discussed in Section 4.1.2.2 [Processing options], page 54. In particular, NoiseChisel uses two tessellations (with everything between them identical except the tile sizes): a fine-grained one with smaller tiles (mainly used in detection) and a more larger tiled one which is used for multi-threaded processing. The common Tessellation options discribed in Section 4.1.2.2 [Processing options], page 54, define all parameters of both tessellations, only the large tile size for the latter tessellation is set through the `--largetilesize` option. To inspect the tessellations on your input dataset, run NoiseChisel with `--checktiles`.

Usage TIP: Frequently use the options starting with `--check`. Depending on what you want to detect in the data, you can often play with the paramters/options for a better result than the default parameters. You can start with `--checkdetection` and `--checksegmentation` for the main steps. For their full list please run:

```
$ astnoisechisel --help | grep check
```

In the sections below, NoiseChisel's options are classified into three general classes to help in easy navigation. Section 7.2.2.1 [General NoiseChisel options], page 166, mainly discusses the options relating to input and those that are shared in both detection and segmentation. Options to configure the detection are described in Section 7.2.2.2 [Detection options], page 169, and Section 7.2.2.3 [Segmentation options], page 173, we discuss how you can fine-tune the segmentation of the detections. Finally in Section 7.2.2.4 [NoiseChisel output], page 175, the format of NoiseChisel's output is discussed. The order of options here follow the same logical order that the respective action takes place within NoiseChisel (note that the output of `--help` is sorted alphabetically).

7.2.2.1 General NoiseChisel options

The options discussed in this section are mainly regarding the input(s), output, and some general processing options that are shared between both detection and segmentation. Recall that you can always see the full list of Gnuastro's options with the `--help` option.

`-k STR`
`--kernel=STR`

File name of kernel to smooth the image before applying the threshold, see Section 6.3.4 [Convolution kernel], page 135. The first step of NoiseChisel is

to convolve/smooth the image and use the convolved image in multiple steps during the processing. It will be used to define (and later apply) the quantile threshold (see `--qthresh`). The convolved image is also used to define the clumps (see Section 3.2.1 and Figure 8 of Akhlaghi and Ichikawa [2015] (`https://arxiv.org/abs/1505.01664`)).

The `--kernel` option is not mandatory. If no kernel is provided, a 2D Gaussian profile with a FWHM of 2 pixels truncated at 5 times the FWHM is used. This choice of the default kernel is discussed in Section 3.1.1 of Akhlaghi and Ichikawa [2015].

`--khdu=STR`
: HDU containing the kernel in the file given to the `--kernel` option.

`-E`
`--skysubtracted`
: If this option is called, it is assumed that the image has already been sky subtracted once. Knowing if the sky has already been subtracted once or not is very important in estimating the Signal to noise ratio of the detections and clumps. In short an extra σ^2_{sky} must be added in the error (noise or denominator in the Signal to noise ratio) for every flux value that is present in the calculation. This can be interpreted as the error in measuring that sky value when it was subtracted by any other program. See Section 3.3 in Akhlaghi and Ichikawa [2015] (`https://arxiv.org/abs/1505.01664`)) for a complete explanation.

`-B FLT`
`--minskyfrac=FLT`
: Minimum fraction (value between 0 and 1) of sky (undetected) areas in a tile for it to be considered in measuring the following detection and segmentation properties.

- Measuring the Signal to noise ratio of false detections during the false detection removal on small tiles.
- Measuring the sky value (average of undetected pixels) on small tiles. Both before the removal of false detections and after it.
- Clump Signal to noise ratio in the sky regions of large files.

Because of the PSF and their intrinsic amorphous properties, astronomical objects (except cosmic rays) never have a clear cutoff and commonly sink into the noise very slowly. Even below the very low thresholds used by NoiseChisel. So when a large fraction of the area of one mesh is covered by detections, it is very plausible that their faint wings are present in the undetected regions (hence causing a bias in any measurement). To get an accurate measurement of the above parameters over the tessellation, tiles that harbor too many detected regions should be excluded. The used tiles are visible in the respective `--check` option of the given step.

`--minnumfalse=INT`
: The minimum number of 'pseudo-detections' (to identify false initial detections) or clumps (to identifying false clumps) found over the un-detected regions to identify a Signal-to-Noise ratio threshold.

The Signal to noise ratio (S/N) of false pseudo-detections and clumps in each tile is found using the quantile of the S/N distribution of the psudo-detections and clumps over the undetected pixels in each mesh. If the number of S/N measurements is not large enough, the quantile will not be accurate (can have large scatter). For example if you set `--detquant=0.99` (or the top 1 percent), then it is best to have at least 100 S/N measurements.

`-L INT[,INT]`
`--largetilesize=INT[,INT]`

The size of each tile for the tessellation with the larger tile sizes. Except for the tile size, all the other parameters for this tessellation are taken from the common options described in Section 4.1.2.2 [Processing options], page 54. The format is identical to that of the `--tilesize` option that is discussed in that section.

`--onlydetection`

If this option is called, no segmentation will be done and the output will only have four extensions (no clumps extension, see Section 7.2.2.4 [NoiseChisel output], page 175). The second extension of the output is not going to be objects but raw detections (a large region will be given one label): labeling is only done based on connectivity. The last two extensions of the output will be the Sky and its Standard deviation.

This option can result in faster processing when only the noise properties of the image are desired for a catalog using another image's labels for example. A common case is when you want to measure colors or SEDs in several images. Let's say you have images in two colors: A and B. For simplicity also assume that they are exactly on the same position in the sky with the same pixel scale.

You choose to set A as a reference, so you run the NoiseChisel fully on A. Then you run NoiseChisel on B with `--onlydetection` since you only need the noise properties of B (for the signal to noise column in its catalog). You can then run MakeCatalog on A normally, see Section 7.3 [MakeCatalog], page 175. To run MakeCatalog on B, you simply set the object and clump labels images to those that NoiseChisel produced for A, see Section 7.3.5 [Invoking MakeCatalog], page 185.

`--grownclumps`

In the output (see Section 7.2.2.4 [NoiseChisel output], page 175) store the grown clumps (or full detected region if only one clump was present in that detection). By default the original clumps are stored as the third extension of the output, if this option is called, it is replaced with the grown clump labels.

`--continueaftercheck`

Continue NoiseChisel after any of the options starting with `--check`. NoiseChisel involves many steps and as a result, there are many checks, allowing to inspect the status of the processing. The results of each step affect the next steps of processing, so, when you are want to check the status of the processing at one step, the time spent to complete NoiseChisel is just wasted/distracting time.

To encourage easier experimentation with the option values, when you use any of the NoiseChisel options that start with `--check`, NoiseChisel will abort once all the desired check file(s) is (are) completed. If you call the `--continueaftercheck` option, you can disable this behavior and ask NoiseChisel to continue with the rest of the processing after completeting the check file(s).

7.2.2.2 Detection options

Detection is the process of separating the pixels in the image into two groups: 1) Signal and 2) Noise. Through the parameters below, you can customize the detection process in NoiseChisel. Recall that you can always see the full list of Gnuastro's options with the `--help` option.

`-r FLT`
`--mirrordist=FLT`
: Maximum distance (as a multiple of error) to estimate the difference between the input and mirror distributions in finding the mode, see Appendix C of Akhlaghi and Ichikawa 2015 (`https://arxiv.org/abs/1505.01664`), also see Section 7.1.3.3 [Quantifying signal in a tile], page 153.

`-Q FLT`
`--modmedqdiff=FLT`
: The maximum acceptable distance between the mode and median, see Section 7.1.3.3 [Quantifying signal in a tile], page 153. The quantile threshold will be found on tiles that satisfy this mode and median difference.

`-t FLT`
`--qthresh=FLT`
: The quantile threshold to apply to the convolved image. The detection process begins with applying a quantile threshold to each of the tiles in the small tessellation. The quantile is only calculated for tiles that don't have any significant signal within them, see Section 7.1.3.3 [Quantifying signal in a tile], page 153. Interpolation is then used to give a value to the un-successful tiles and it is finally smoothed.

  The quantile value is a floating point value between 0 and 1. Assume that we have sorted the N data elements of a distribution (the pixels in each mesh on the convolved image). The quantile (q) of this distribution is the value of the element with an index of (the nearest integer to) $q \times N$ in the sorted data set. After thresholding is complete, we will have a binary (two valued) image. The pixels above the threshold are known as foreground pixels (have a value of 1) while those which lie below the threshold are known as background (have a value of 0).

`--smoothwidth=INT`
: Width of flat kernel used to smooth the interpolated quantile thresholds, see `--qthresh` for more.

`--checkqthresh`
: Check the quantile threshold values on the mesh grid. A file suffixed with `_qthresh.fits` will be created showing each step. With this option, NoiseChisel

will abort as soon as quantile estimation has been completed, allowing you to inspect the steps leading to the final quantile threshold, this can be disabled with `--continueaftercheck`. By default the output will have the same pixel size as the input, but with the `--oneelempertile` option, only one pixel will be used for each tile (see Section 4.1.2.2 [Processing options], page 54).

`-e INT`
`--erode=INT`

The number of erosions to apply to the binary thresholded image. Erosion is simply the process of flipping (from 1 to 0) any of the foreground pixels that neighbor a background pixel. In a 2D image, there are two kinds of neighbors, 4-connected and 8-connected neighbors. You can specify which type of neighbors should be used for erosion with the `--erodengb` option, see below.

Erosion has the effect of shrinking the foreground pixels. To put it another way, it expands the holes. This is a founding principle in NoiseChisel: it exploits the fact that with very low thresholds, the holes in the very low surface brightness regions of an image will be smaller than regions that have no signal. Therefore by expanding those holes, we are able to separate the regions harboring signal.

`--erodengb=INT`

The type of neighborhood (structuring element) used in erosion, see `--erode` for an explanation on erosion. Only two integer values are acceptable: 4 or 8. In 4-connectivity, the neighbors of a pixel are defined as the four pixels on the top, bottom, right and left of a pixel that share an edge with it. The 8-connected neighbors on the other hand include the 4-connected neighbors along with the other 4 pixels that share a corner with this pixel. See Figure 6 (a) and (b) in Akhlaghi and Ichikawa (2015) for a demonstration.

`--noerodequant`

Pure erosion is going to carve off sharp and small objects completely out of the detected regions. This option can be used to avoid missing such sharp and small objects (which have significant pixels, but not over a large area). All pixels with a value larger than the significance level specified by this option will not be eroded during the erosion step above. However, they will undergo the erosion and dilation of the opening step below.

Like the `--qthresh` option, the significance level is determined using the quantile (a value between 0 and 1). Just as a reminder, in the normal distribution, 1σ, 1.5σ, and 2σ are approximately on the 0.84, 0.93, and 0.98 quantiles.

`-p INT`
`--opening=INT`

Depth of opening to be applied to the eroded binary image. Opening is a composite operation. When opening a binary image with a depth of n, n erosions (explained in `--erode`) are followed by n dilations. Simply put, dilation is the inverse of erosion. When dilating an image any background pixel is flipped (from 0 to 1) to become a foreground pixel. Dilation has the effect of fattening the foreground. Note that in NoiseChisel, the erosion which is part of opening is independent of the initial erosion that is done on the thresholded image (explained in `--erode`). The structuring element for the opening can be specified

with the `--openingngb` option. Opening has the effect of removing the thin foreground connections (mostly noise) between separate foreground 'islands' (detections) thereby completely isolating them. Once opening is complete, we have *initial* detections.

`--openingngb=INT`
: The structuring element used for opening, see `--erodengb` for more information about a structuring element.

`-s FLT,FLT`
`--sigmaclip=FLT,FLT`
: The σ-clipping parameters, see Section 7.1.2 [Sigma clipping], page 149. This option takes two values which are separated by a comma (`,`). Each value can either be written as a single number or as a fraction of two numbers (for example `3,1/10`). The first value to this option is the multiple of σ that will be clipped (α in that section). The second value is the exit criteria. If it is less than 1, then it is interpretted as tolerance and if it is larger than one it is assumed to be the fixed number of iterations. Hence, in the latter case the value must be an integer.

`--checkdetsky`
: Check the initial approximation of the sky value and its standard deviation in a FITS file ending with `_detsky.fits`. With this option, NoiseChisel will abort as soon as the sky value used for defining pseudo-detections is complete. This allows you to inspect the steps leading to the final quantile threshold, this behavior can be disabled with `--continueaftercheck`. By default the output will have the same pixel size as the input, but with the `--oneelempertile` option, only one pixel will be used for each tile (see Section 4.1.2.2 [Processing options], page 54).

`-R FLT`
`--dthresh=FLT`
: The detection threshold: a multiple of the initial sky standard deviation added with the initial sky approximation (which you can inspect with `--checkdetsky`). This flux threshold is applied to the initially undetected regions on the unconvolved image. The background pixels that are completely engulfed in a 4-connected foreground region are converted to background (holes are filled) and one opening (depth of 1) is applied over both the initially detected and undetected regions. The Signal to noise ratio of the resulting 'psudo-detections' are used to identify true vs. false detections. See Section 3.1.5 and Figure 7 in Akhlaghi and Ichikawa (2015) for a very complete explanation.

`-i INT`
`--detsnminarea=INT`
: The minimum area to calculate the Signal to noise ratio on the psudo-detections of both the initially detected and undetected regions. When the area in a psudo-detection is too small, the Signal to noise ratio measurements will not be accurate and their distribution will be heavily skewed to the positive. So

it is best to ignore any psudo-detection that is smaller than this area. Use `--detsnhistnbins` to check if this value is reasonable or not.

`--checkdetsn`
: Save the S/N values of the pseudo-detections and dilated detections into three files ending with `_detsn_sky.XXX`, `_detsn_det.XXX`, and `_detsn_dilated.XXX`. The `.XXX` is determined from the `--tableformat` option (see Section 4.1.2.1 [Input/Output options], page 52, for example `.txt` or `.fits`). You can use these to inspect the S/N values and their distribution (in combination with the `--checkdetection` option to see where the pseudo-detections are). You can use Gnuastro's Section 7.1 [Statistics], page 148, to make a histogram of the distribution or any other analysis you would like for better understanding of the distribution (for example through a histogram).

  With this option, NoiseChisel will abort as soon as the tables are created. This allows you to inspect the steps leading to the final quantile threshold, this behavior (to abort NoiseChisel) can be disabled with `--continueaftercheck`.

`-c FLT`
`--detquant=FLT`
: The quantile of the Signal to noise ratio distribution of the psudo-detections in each mesh to use for filling the large mesh grid. Note that this is only calculated for the large mesh grids that satisfy the minimum fraction of undetected pixels (value of `--minbfrac`) and minimum number of psudo-detections (value of `--minnumfalse`).

`-d INT`
`--dilate=INT`
: Number of times to dilate the final true detections. See the explanations in `--opening` for more information on dilation. The structuring element for this final dilation is fixed to an 8-connected neighborhood. This is because astronomical objects, except cosmic rays, never have a clear cutoff, so all the 8-pixels connected to the border pixels of a detection might harbor data.

`--checkdetection`
: Every step of the detection process will be added as an extension to a file with the suffix `_det.fits`. Going through each would just be a repeat of the explanations above and also of those in Akhlaghi and Ichikawa (2015). The extension label should be sufficient to recognize which step you are observing. Viewing all the steps can be the best guide in choosing the best set of parameters. With this option, NoiseChisel will abort as soon as a snapshot of all the detection process is saved. This behavior can be disabled with `--continueaftercheck`.

`--checksky`
: Check the derivation of the final sky and its standard deviation values on the mesh grid. With this option, NoiseChisel will abort as soon as the sky value is estimated over the image (on each tile). This behavior can be disabled with `--continueaftercheck`. By default the output will have the same pixel size as the input, but with the `--oneelempertile` option, only one pixel will be used for each tile (see Section 4.1.2.2 [Processing options], page 54).

7.2.2.3 Segmentation options

Segmentation is the process of (possibly) breaking up a detection into multiple segments (technically called *objects* and *clumps* in NoiseChisel). In deep surveys segmentation becomes particularly important because we will be detecting more diffuse flux so galaxy images are going to overlap more. It is thus very important to be able separate the pixels within a detection.

In NoiseChisel, segmentation is done by first finding the 'true' clumps over a detection and then expanding those clumps to a certain flux limit. True clumps are found in a process very similar to the true detections explained in Section 7.2.2.2 [Detection options], page 169, see Akhlaghi and Ichikawa [2015] (`https://arxiv.org/abs/1505.01664`) for more information. If the connections between the grown clumps are weaker than a given threshold, the grown clumps are considered to be separate objects.

`-m INT`
`--segsnminarea=INT`
: The minimum area which a clump in the undetected regions should have in order to be considered in the clump Signal to noise ratio measurement. If this size is set to a small value, the Signal to noise ratio of false clumps will not be accurately found. It is recommended that this value be larger than the value to `--detsnminarea`. Because the clumps are found on the convolved (smoothed) image while the psudo-detections are found on the input image. You can use `--checkclumpsn` and `--checksegmentation` to see if your chosen value is reasonable or not.

`--checkclumpsn`
: Save the S/N values of the clumps into two files ending with `_clumpsn_sky.XXX` and `_clumpsn_det.XXX`. The `.XXX` is determined from the `--tableformat` option (see Section 4.1.2.1 [Input/Output options], page 52, for example `.txt` or `.fits`). You can use these to inspect the S/N values and their distribution (in combination with the `--checksegmentation` option to see where the clumps are). You can use Gnuastro's Section 7.1 [Statistics], page 148, to make a histogram of the distribution (ready for plotting in a text file, or a crude ASCII-art demonstration on the command-line).

  With this option, NoiseChisel will abort as soon as the two tables are created. This allows you to inspect the steps leading to the final S/N quantile threshold, this behavior can be disabled with `--continueaftercheck`.

`-g FLT`
`--segquant=FLT`
: The quantile of the noise clump Signal to noise ratio distribution. This value is used to identify true clumps over the detected regions. You can get the full distribution of clumps S/Ns over the undetected areas with the `--checkclumpsn` option and see them with `--checksegmentation`.

`-v`
`--keepmaxnearriver`
: Keep a clump whose maximum flux is 8-connected to a river pixel. By default such clumps over detections are considered to be noise and are removed irrespective of their brightness (see Section 8.1.3 [Flux Brightness and magnitude],

page 201). Over large profiles, that sink into the noise very slowly, noise can cause part of the profile (which was flat without noise) to become a very large and with a very high Signal to noise ratio. In such cases, the pixel with the maximum flux in the clump will be immediately touching a river pixel.

`-G FLT`
`--gthresh=FLT`

Threshold (multiple of the sky standard deviation added with the sky) to stop growing true clumps. Once true clumps are found, they are set as the basis to segment the detected region. They are grown until the threshold specified by this option.

`-y INT`
`--minriverlength=INT`

The minimum length of a river between two grown clumps for it to be considered in Signal to noise ratio estimations. Similar to `--segsnminarea` and `--detsnminarea`, if the length of the river is too short, the Signal to noise ratio can be noisy and unreliable. Any existing rivers shorter than this length will be considered as non-existent, independent of their Signal to noise ratio. Since the clumps are grown on the input image, this value should best be similar to the value of `--detsnminarea`. Recall that the clumps were defined on the convolved image so `--segsnminarea` was larger than `--detsnminarea`.

`-O FLT`
`--objbordersn=FLT`

The maximum Signal to noise ratio of the rivers between two grown clumps in order to consider them as separate 'objects'. If the Signal to noise ratio of the river between two grown clumps is larger than this value, they are defined to be part of one 'object'. Note that the physical reality of these 'objects' can never be established with one image, or even multiple images from one broad-band filter. Any method we devise to define 'object's over a detected region is ultimately subjective.

Two very distant galaxies or satellites in one halo might lie in the same line of sight and be detected as clumps on one detection. On the other hand, the connection (through a spiral arm or tidal tail for example) between two parts of one galaxy might have such a low surface brightness that they are broken up into multiple detections or objects. In fact if you have noticed, exactly for this purpose, this is the only Signal to noise ratio that the user gives into NoiseChisel. The 'true' detections and clumps can be objectively identified from the noise characteristics of the image, so you don't have to give any hand input Signal to noise ratio.

`--checksegmentation`

A file with the suffix `_seg.fits` will be created. This file keeps all the relevant steps in finding true clumps and segmenting the detections into multiple objects in various extensions. Having read the paper or the steps above. Examining this file can be an excellent guide in choosing the best set of parameters. Note that calling this function will significantly slow NoiseChisel. In verbose mode (without the `--quiet` option, see Section 4.1.2.3 [Operating mode options],

page 55) the important steps (along with their extension names) will also be reported.

With this option, NoiseChisel will abort as soon as the two tables are created. This behavior can be disabled with `--continueaftercheck`.

7.2.2.4 NoiseChisel output

The default name and directory of the outputs are explained in Section 4.8 [Automatic output], page 77. NoiseChisel's default output (when none of the options starting with `--check` or the `--output` option are called) is one file ending with `_labeled.fits`. This file has the extensions listed below:

1. A copy of the input image, a copy is placed here for the following reasons:
   - By flipping through the extensions, a user can check how accurate the detection and segmentation process was.
   - All the inputs to MakeCatalog (see Section 7.3 [MakeCatalog], page 175) are included in this one file which makes the running of MakeCatalog after NoiseChisel very easy.
2. The object/detection labels. Each pixel in the input image is given a label in this extension, the labels start from one. If the `--onlydetection` option is given, each large connected part of the image has one label. Without that option, this extension is going to show the labels of the objects that are found after segmentation. The total number of labels is stored as the value to the `NOBJS/NDETS` keyword in the header of this extension. This number is also printed in verbose mode.
3. The clump labels when `--onlydetection` is not called. All the pixels in the input image that belong to a true clump are given a positive label in this extension. The detected regions that were not a clump are given a negative value to clearly identify the sky noise from the diffuse detections. The total number of clumps in this image is stored in the `NCLUMPS` keyword of this extension and printed in verbose output.

   If the `--grownclumps` option is called, or a value of `1` is given to it in any of the configuration files, then instead of the original clump regions, the grown clumps will be stored in this extension. Note that if there is only one clump (or no clumps) over a detected region, then the whole detected region is given a label of 1.
4. The final sky value on each pixel. See Section 7.1.3 [Sky value], page 150, for a complete explanation.
5. Similar to the previous mesh but for the standard deviation on each pixel.

7.3 MakeCatalog

At the lowest level, a dataset (for example an image) is just a collection of values, placed after each other in any number of dimensions (for example an image is a 2D dataset). Each data-element (pixel) just has two properties: its position (relative to the rest) and its value. The entire input dataset (a large image for example) is rarely treated as a singular entity for higher-level analysis[9]. You want to know the properties of the scientifically intresting targets

---

[9] In low-level reduction/preparation of a dataset, you do infact treat a whole image as a single entity. You can derive the over-all properties of all the pixels in a dataset with Gnuastro's Statistics program (see Section 7.1 [Statistics], page 148)

that are embedded in it. For example the magnitudes, positions and elliptical properties of the galaxies that are in the image. MakeCatalog is Gnuastro's program to derive higher-level information for *pre-defined* regions of a dataset. The role of MakeCatalog in a scientific analysis and the benefits of this model of data-analysis (were detection/identification is separated from measurement) is discussed in Akhlaghi [2016] (`https://arxiv.org/abs/1611.06387v1`). We strongly recommend reading this short paper for a better understanding of this methodology and use MakeCatalog most effectively. However, that paper cannot undergo any more change, so this manual is the definitive guide.

As discussed above, you have to define the regions of a dataset that you are interested in *before* running MakeCatalog. MakeProfiles currently uses labeled dataset(s) for this job. A labeled dataset for a given input dataset has the size/dimensions as the input, but its pixels have an integer type (see Section 4.4 [Numeric data types], page 64)[10]: all pixels with the same label (integers larger and equal to one) are used to generate the requested output columns of MakeCatalog for the row of their labeled value. For example, the flux weighted average position of all the pixels with a label of 42 will be considered as the central position[11] of the 42nd row of the output catalog. Pixels with labels equal to or smaller than zero will be ignored by MakeCatalog. In other words, the number of rows of the output catalog will be determined from the labeled image.

The labeled image maybe created with any tool[12]. Within Gnuastro you can use these two solutions depending on a-priori/parametric knowledge of the targets you want to study:

- When you already know the positions and parametric (for example circular or elliptical) properties of the targets, you can use Section 8.1 [MakeProfiles], page 195, to generate a labeled image from another catalog. This is also known as aperture photometry (the apertures are defined a-priori).
- When the shape of your targets cannot be parametrized accurately (for example galaxies), or if you don't know the number/shape of the targets in the image, you can use Gnuastro's NoiseChisel program to detect and segment (make labeled images of) the *objects* and *clumps* in the input image, see Section 7.2 [NoiseChisel], page 163.

7.3.1 Detection and catalog production

As discussed above (Section 7.3 [MakeCatalog], page 175), NoiseChisel (Gnuastro's signal detection tool, see Section 7.2 [NoiseChisel], page 163) does not produce any catalog of the detected objects. However, most other common tools in astronomical data-analysis (for example SExtractor[13]) merge the two processes into one. Gnuastro's modularized methodology is therefore new to many experienced astronomers and deserves a short review here. Further discussion on the benefits of this methodology can be seen in Akhlaghi [2016] (`https://arxiv.org/abs/1611.06387v1`).

---

[10] If the program you used to generate the labeled image only outputs floating point types, but you know it only has integer valued pixels that are stored in a floating point container, you can use Gnuastro's Arithmetic program (see Section 6.2 [Arithmetic], page 108) to change the numerical data type of the image (`flabel.fits`) to an integer type image (`label.fits`) with a command like below:

```
$ astarithmetic flabel.fits int32 --output=label.fits
```

[11] See Section 7.3.3 [Measuring elliptical parameters], page 181, for a discussion on this and the derivation of positional parameters.

[12] For example, you can even use a graphic user interfase image editing tool like the GNU Image Manipulation Program (or GIMP) and use Gnuastro's ConvertType to convert it to a FITS file.

[13] `https://www.astromatic.net/software/sextractor`

To simplify catalog production from a raw input image in Gnuastro, NoiseChisel's output (see Section 7.2.2.4 [NoiseChisel output], page 175) can be directly fed into MakeCatalog. This is good when no further customization is necessary and you want a fast/simple. But the modular approach taken by Gnuastro has many benefits that will become more apparent as you get more experienced in astronomical data analysis and want to be more creative in using your valuable data for the exciting scientific project you are working on. In short the reasons for this modularity can be classified as below:

- Complexity of a monolith: Adding in a catalog functionality to the detector program will add several more steps (and many more options) to its processing that can equally well be done outside of it. This makes following what the program does harder for the users and developers, it can also potentially add many bugs.

  As an example, if the parameter you want to measure over one profile is not provided by the developers of MakeCatalog. You can simply open this tiny little program and add your desired calculation easily. This process is discussed in Section 7.3.4 [Adding new columns to MakeCatalog], page 184. However, if making a catalog was part of NoiseChisel, it would require a lot of energy to understand all the steps and internal structures of that large program (the most complex in Gnuastro) in order to add desired parameter in a catalog.

- Simplicity/robustness of independent, modular tools: making a catalog is a logically separate process from labeling (detection and segmentation). A user might want to do certain operations on the labeled regions before creating a catalog for them. Another user might want the properties of the same pixels/objects in another image (another filter for example) to measure the colors or SED fittings.

  Here is an example of doing both: suppose you have images in various broad band filters at various resolutions and orientations. The image of one color will thus not lie exactly on another or even be in the same scale. However, it is imperative that the same pixels be used in measuring the colors of galaxies.

  To solve the problem, NoiseChisel can be run on the reference image to generate the labeled image. After wards, the labeled image can be warped into the grid of the other color (using Section 6.4 [Warp], page 138). MakeCatalog will then generate the same catalog for both colors (with the different labeled images). It is currently customary to warp the images to the same pixel grid, however, modification of the scientific dataset is very harmful for the data and creates correlated noise. It is much more accurate to do the transformations on the labeled image.

7.3.2 Quantifying measurement limits

No measurement on a real dataset can be perfect: you can only reach a certain level/limit of accuracy. Therefore, a meaningful (scientific) analysis requires an understanding of these limits for the dataset and your analysis tools: different datasets (images in the case of MakeCatalog) have different noise properties and different detection methods (one method/algorith/software that is run with a different set of parameters is considered as a different detection method) will have different abilities to detect or measure certain kinds of signal (astronomical objects) and their properties in an image. Hence, quantifying the detection and measurement limitations with a particular dataset and analysis tool is the most crucial/critical aspect of any high-level analysis.

In this section we discuss some of the most general limits that are very important in any astronomical data analysis and how MakeCatalog makes it easy to find them. Depending on the higher-level analysis, there are more tests that must be done, but these are usually necessary in any case. In astronomy, it is common to use the magnitude (a unit-less scale) and physical units, see Section 8.1.3 [Flux Brightness and magnitude], page 201. Therefore all the measurements discussed here are defined in units of magnitudes.

Surface brightness limit (of whole dataset)

As we make more observations on one region of the sky, and add the observations into one dataset, we are able to decrease the standard deviation of the noise in each pixel[14]. Qualitatively, this decrease manifests its self by making fainter (per pixel) parts of the objects in the image more visible. Technically, this is known as surface brightness. Quantitatively, it increases the Signal to noise ratio, since the signal increases faster than noise with more data. It is very important to have in mind that here, noise is defined per pixel (or in the units of our data measurement), not per object.

You can think of the noise as muddy water that is completely covering a flat ground[15] with some regions higher than the others[16] in it. In this analogy, height (from the ground) is *surface brightness*. Let's assume that in your first observation the muddy water has just been stirred and you can't see anything through it. As you wait and make more observations, the mud settles down and the *depth* of the transparent water increases, making the summits of hills visible. As the depth of clear water increases, the parts of the hills with lower heights (less parts with lower surface brightness) can be seen more clearly.

The outputs of NoiseChisel include the Sky standard deviation (σ) on every group of pixels (a mesh) that were calculated from the undetected pixels in that mesh, see Section 4.6 [Tessellation], page 72, and Section 7.2.2.4 [NoiseChisel output], page 175. Let's take σ_m as the median σ over the successful meshes in the image (prior to interpolation or smoothing).

On different instruments pixels have different physical sizes (for example in micro-meters, or spatial angle over the sky), nevertheless, a pixel is our unit of data collection. In other words, while quantifying the noise, the physical or projected size of the pixels is irrelevant. We thus define the Surface brightness limit or *depth*, in units of magnitude/pixel, of a data-set, with zeropoint magnitude z, with the nth multiple of σ_m as (see Section 8.1.3 [Flux Brightness and magnitude], page 201):

$$SB_{\rm Pixel} = -2.5 \times \log_{10}(n\sigma_m) + z$$

As an example, the XDF survey covers part of the sky that the Hubble space telescope has observed the most (for 85 orbits) and is consequently very small

[14] This is true for any noisy data, not just astronomical images

[15] The ground is the sky value in this analogy, see Section 7.1.3 [Sky value], page 150. Note that this analogy only holds for a flat sky value across the surface of the image or ground.

[16] The peaks are the brightest parts of astronomical objects in this analogy.

(~ 4 arcmin^2). On the other hand, the CANDELS survey, is one of the widest multi-color surveys covering several fields (about 720 arcmin^2) but its deepest fields have only 9 orbits observation. The depth of the XDF and CANDELS-deep surveys in the near infrared WFC3/F160W filter are respectively 34.40 and 32.45 magnitudes/pixel. In a single orbit image, this same field has a depth of 31.32. Recall that a larger magnitude corresponds to less brightness.

The low-level magnitude/pixel measurement above is only useful when all the datasets you want to use belong to one instrument (telescope and camera). However, you will often find yourself using datasets from various instruments with different pixel scales (projected pixel sizes). If we know the pixel scale, we can obtain a more easily comparable surface brightness limit in units of: magnitude/arcsec^2. Let's assume that the dataset has a zeropoint value of z, and every pixel is p arcsec^2 (so A/p is the number of pixels that cover an area of A arcsec^2). If the nth multiple of σ_m is desired, then the surface brightness (in units of magnitudes per A arcsec^2) is[17]:

$$SB_{\rm Projected} = -2.5 \times \log_{10}\left(n\sigma_m\sqrt{\frac{A}{p}}\right) + z$$

Note that this is an extrapolation of the actually measured value of σ_m (which was per pixel). So it should be used with extreme care (for example the dataset must have an approximately flat depth). For each detection over the dataset, you can estimate an upper-limit magnitude which actually uses the detection's area/footprint. It doesn't extrapolate and even accounts for correlated noise features. Therefore, the upper-limit magnitude is a much better measure of your dataset's surface brightness limit for each particular object.

MakeCatalog will calculate the input dataset's $SB_{\rm Pixel}$ and $SB_{\rm Projected}$ and write them as comments/meta-data in the output catalog(s). Just note that $SB_{\rm Projected}$ is only calculated if the input has World Coordinate System (WCS).

Completeness limit (of each detection)

As the surface brightness of the objects decreases, the ability to detect them will also decrease. An important statistic is thus the fraction of objects of similar morphology and brightness that will be identified with our detection algorithm/parameters in the given image. This fraction is known as completeness. For brighter objects, completeness is 1: all bright objects that might exist over the image will be detected. However, as we go to lower surface brightness objects, we fail to detect some and gradually we are not able to detect anything any more. For a given profile, the magnitude where the completeness drops below a certain level usually above 90% is known as the completeness limit.

Another important parameter in measuring completeness is purity: the fraction of true detections to all true detections. In effect purity is the measure of contamination by false detections: the higher the purity, the lower the contamination. Completeness and purity are anti-correlated: if we can allow a large

[17] If we have N datasets, each with noise σ, the noise of a combined dataset will increase as $\sqrt{N}\sigma$.

number of false detections (that we might be able to remove by other means), we can significantly increase the completeness limit.

One traditional way to measure the completeness and purity of a given sample is by embedding mock profiles in regions of the image with no detection. However in such a study we must be really careful to choose model profiles as similar to the target of interest as possible.

Magnitude measurement error (of each detection)

Any measurement has an error and this includes the derived magnitude for an object. Note that this value is only meaningful when the object's magnitude is brighter than the upper-limit magnitude (see the next items in this list). As discussed in Section 8.1.3 [Flux Brightness and magnitude], page 201, the magnitude (M) of an object with brightness B and Zeropoint magnitude z can be written as:

$$M = -2.5\log_{10}(B) + z$$

Calculating the derivative with respect to B, we get:

$$\frac{dM}{dB} = \frac{-2.5}{B \times ln(10)}$$

From the Tailor series ($\Delta M = dM/dB \times \Delta B$), we can write:

$$\Delta M = \left|\frac{-2.5}{ln(10)}\right| \times \frac{\Delta B}{B}$$

But, $\Delta B/B$ is just the inverse of the Signal-to-noise ratio (S/N), so we can write the error in magnitude in terms of the signal-to-noise ratio:

$$\Delta M = \frac{2.5}{S/N \times ln(10)}$$

MakeCatalog uses this relation to estimate the magnitude errors. The signal-to-noise ratio is calculated in different ways for clumps and objects (see Akhlaghi and Ichikawa [2015] (`https://arxiv.org/abs/1505.01664`)), but this single equation can be used to estimate the measured magnitude error afterwards for any type of target.

Upper limit magnitude (of each detection)

Due to the noisy nature of data, it is possible to get arbitrarily low values for a faint object's brightness (or arbitrarily high magnitudes). Given the scatter caused by the noise, such small values are meaningless: another similar depth observation will give a radically different value. This problem is most common when you use one image/filter to generate target labels (which specify which pixels belong to which object, see Section 7.2.2.4 [NoiseChisel output], page 175,

and Section 7.3 [MakeCatalog], page 175) and another image/filter to generate a catalog for measuring colors.

The object might not be visible in the filter used for the latter image, or the image *depth* (see above) might be much shallower. So you will get unreasonably faint magnitudes. For example when the depth of the image is 32 magnitudes, a measurement that gives a magnitude of 36 for a ~ 100 pixel object is clearly unreliable. In another similar depth image, we might measure a magnitude of 30 for it, and yet another might give 33. Furthermore, due to the noise scatter so close to the depth of the data-set, the total brightness might actually get measured as a negative value, so no magnitude can be defined (recall that a magnitude is a base-10 logarithm).

Using such unreliable measurements will directly affect our analysis, so we must not use them. However, all is not lost! Given our limited depth, there is one thing we can deduce about the object's magnitude: we can say that if something actually exists here (possibly buried deep under the noise), it must have a magnitude that is fainter than an *upper limit magnitude*. To find this upper limit magnitude, we place the object's footprint (segmentation map) over random parts of the image where there are no detections, so we only have pure (possibly correlated) noise and undetected objects. Doing this a large number of times will give us a distribution of brightness values. The standard deviation (σ) of that distribution can be used to quantify the upper limit magnitude.

Traditionally, faint/small object photometry was done using fixed circular apertures (for example with a diameter of N arc-seconds). In this way, the upper limit was like the depth discussed above: one value for the whole image. But with the much more advanced hardware and software of today, we can make customized segmentation maps for each object. The number of pixels (are of the object) used directly affects the final distribution and thus magnitude. Also the image correlated noise might actually create certain patters, so the shape of the object can also affect the result. So in MakeCatalog, the upper limit magnitude is found for each object in the image separately. Not one value for the whole image.

7.3.3 Measuring elliptical parameters

The shape or morphology of a target is one of the most commonly desired paramters of a target. Here, we will review the derivation of the most basic/simple morphological parameters are estimated: the elliptical parameters for a set of labeled pixels. The elliptical parameters are: the (semi-)major axis, the (semi-)minor axis and the position angle along with the central position of the profile. The derivations below follow the SExtractor manual derivations with some added explanations for easier reading.

Let's begin with one dimension for simplicity: Assume we have a set of N values B_i (keeping the spatial distribution of brightness for example), each at position x_i. The simplest parameter we can define is the geometric center of the object (x_g) (ignoring the brightness values): $x_g = (\sum_i x_i)/N$. *Moments* are defined to incorporate both the value (brightness) and position of the data. The first moment can be written as:

$$\overline{x} = \frac{\sum_i B_i x_i}{\sum_i B_i}$$

This is essentially the weighted (by B_i) mean position. The geometric center (x_g, defined above) is a special case of this with all $B_i = 1$. The second moment is essentially the variance of the distribution:

$$\overline{x^2} \equiv \frac{\sum_i B_i (x_i - \overline{x})^2}{\sum_i B_i} = \frac{\sum_i B_i x_i^2}{\sum_i B_i} - 2\overline{x}\frac{\sum_i B_i x_i}{\sum_i B_i} + \overline{x}^2 = \frac{\sum_i B_i x_i^2}{\sum_i B_i} - \overline{x}^2$$

The last step was done from the definition of $\overline{x}$. Hence, the square root of $\overline{x^2}$ is the spatial standard deviation (along the one-dimensional) of this particular brightness distribution (B_i). Crudely (or qualitatively), you can think of its square root as the distance (from $\overline{x}$) which contains a specific amount of the flux (depending on the B_i distribution). Similar to the first moment, the geometric second moment can be found by setting all $B_i = 1$. So while the first moment quantified the position of the brightness distribution, the second moment quantifies how that brightness is dispersed about the first moment. In other words, it quantifies how “sharp” the object’s image is.

Before continuing to two dimensions and the derivation of the elliptical parameters, let’s pause for an important implementation technicality. You can ignore this paragraph if you don’t want to implement these concepts. The basic definition (first fraction for $\overline{x^2}$) can be used without any major problem. However, using this fraction requires two runs over the data: one run to find $\overline{x}$ and one run to find $\overline{x^2}$, this can be slow. However, using the last fraction above, we can estimate both the first and second moments in one run (since the $-\overline{x}^2$ term can easily be added later). The logarithmic nature of floating point number digitization creates a complication in this approach: suppose the object is located between pixels 10000 and 10020. Hence the target’s pixels are only distributed over 20 pixels (with a standard deviation < 20), while the mean has a value of ~ 10000. The $\sum_i B_i^2 x_i^2$ will go to very very large values while the individual pixel differences will be much smaller, this will lower the accuracy of our calculation due to the limited accuracy of floating point operations. The variance only depends on the distance of each point from the mean, so we can shift all position by a constant/arbitrary K which is much closer to the mean: $\overline{x - K} = \overline{x} - K$. Hence we can calculate the second order moment using:

$$\overline{x^2} = \frac{\sum_i B_i (x_i - K)^2}{\sum_i B_i} - (\overline{x} - K)^2$$

The closer K is to $\overline{x}$, the better (the sums of squares will involve smaller numbers), as long as K is within the object limits (in the example above: $10000 \leq K \leq 10020$), the floating point error induced in our calculation will be negligible. For the most simplest implementation, MakeCatalog takes K to be the smallest position of the object in each dimension. Since K is arbitrary and an implementation/technical detail, we will ignore it for the remainder of this discussion.

In two dimensions, the mean and variances can be written as:

$$\overline{x} = \frac{\sum_i B_i x_i}{B_i}, \quad \overline{x^2} = \frac{\sum_i B_i x_i^2}{\sum_i B_i} - \overline{x}^2$$

$$\overline{y} = \frac{\sum_i B_i y_i}{B_i}, \quad \overline{y^2} = \frac{\sum_i B_i y_i^2}{\sum_i B_i} - \overline{y}^2$$

$$\overline{xy} = \frac{\sum_i B_i x_i y_i}{\sum_i B_i} - \overline{x} \times \overline{y}$$

If an elliptical profile's major axis exactly lies along the x axis, then $\overline{x^2}$ will be directly proportional with the profile's major axis, $\overline{y^2}$ with its minor axis and $\overline{xy} = 0$. However, in reality we are not that lucky and (assuming galaxies can be parametrized as an ellipse) the major axis of galaxies can be in any direction on the image (in fact this is one of the core principles behind weak-lensing by shear estimation). So the purpose of the remainder of this section is to define a strategy to measure the position angle and axis ratio of some randomly positioned ellipses in an image, using the raw second moments that we have calculated above in our image coordinates.

Let's assume we have rotated the galaxy by θ, the new second order moments are:

$$\overline{x_\theta^2} = \overline{x^2}\cos^2\theta + \overline{y^2}\sin^2\theta - 2\overline{xy}\cos\theta\sin\theta$$

$$\overline{y_\theta^2} = \overline{x^2}\sin^2\theta + \overline{y^2}\cos^2\theta + 2\overline{xy}\cos\theta\sin\theta$$

$$\overline{xy_\theta} = \overline{x^2}\cos\theta\sin\theta - \overline{y^2}\cos\theta\sin\theta + \overline{xy}(\cos^2\theta - \sin^2\theta)$$

The best θ (θ_0, where major axis lies along the x_θ axis) can be found by:

$$\left.\frac{\partial\overline{x_\theta^2}}{\partial\theta}\right|_{\theta_0} = 0$$

Taking the derivative, we get:

$$2\cos\theta_0\sin\theta_0(\overline{y^2} - \overline{x^2}) + 2(\cos^2\theta_0 - \sin^2\theta_0)\overline{xy} = 0$$

When $\overline{x^2} \neq \overline{y^2}$, we can write:

$$\tan 2\theta_0 = 2\frac{\overline{xy}}{\overline{x^2} - \overline{y^2}}.$$

MakeCatalog uses the standard C math library's `atan2` function to estimate θ_0, which we define as the position angle of the ellipse. To recall, this is the angle of the major axis of the ellipse with the x axis. By definition, when the elliptical profile is rotated by θ_0, then $\overline{xy_{\theta_0}} = 0$, $\overline{x_{\theta_0}^2}$ will be the extent of the maximum variance and $\overline{y_{\theta_0}^2}$ the extent of the minimum variance (which are perpendicular for an ellipse). Replacing θ_0 in the equations above for $\overline{x_\theta}$ and $\overline{y_\theta}$, we can get the semi-major (A) and semi-minor (B) lengths:

$$A^2 \equiv \overline{x^2_{\theta_0}} = \frac{\overline{x^2}+\overline{y^2}}{2} + \sqrt{\left(\frac{\overline{x^2}-\overline{y^2}}{2}\right)^2 + \overline{xy}^2}$$

$$B^2 \equiv \overline{y^2_{\theta_0}} = \frac{\overline{x^2}+\overline{y^2}}{2} - \sqrt{\left(\frac{\overline{x^2}-\overline{y^2}}{2}\right)^2 + \overline{xy}^2}$$

As a summary, it is important to remember that the units of A and B are in pixels (the standard deviation of a positional distribution) and that they represent the spatial light distribution of the object in both image dimensions (rotated by θ_0). When the object cannot be represented as an ellipse, this interpretation breaks down: $\overline{xy_{\theta_0}} \neq 0$ and $\overline{y^2_{\theta_0}}$ will not be the direction of minimum variance.

7.3.4 Adding new columns to MakeCatalog

MakeCatalog is designed to allow easy addition of different measurements over a labeled image (see Akhlaghi [2016] (`https://arxiv.org/abs/1611.06387v1`)). A check-list style description of necessary steps to do that is described in this section. The common development characteristics of MakeCatalog and other Gnuastro programs is explained in Chapter 11 [Developing], page 320. We strongly encourage you to have a look at that chapter to greatly simplify your navigation in the code. After adding and testing your column, you are most welcome (and encouraged) to share it with us so we can add to the next release of Gnuastro for everyone else to also benefit from your efforts.

MakeCatalog will first pass over each label's pixels two times and do necessary raw/internal calculations. Once the passes are done, it will use the raw information for filling the final catalog's columns. In the first pass it will gather mainly object information and in the second run, it will mainly focus on the clumps, or any other measurement that needs an output from the first pass. These two passes are designed to be raw summations: no extra processing. This will allow parallel processing and simplicity/clarity. So if your new calculation, needs new raw information from the pixels, then you will need to also modify the respective `mkcatalog_first_pass` and `mkcatalog_second_pass` functions (both in `bin/mkcatalog/mkcatalog.c`) and define new raw table columns in `main.h` (hopefully the comments in the code are clear enough).

In all these different places, the final columns are sorted in the same order (same order as Section 7.3.5 [Invoking MakeCatalog], page 185). This allows a particular column/option to be easily found in all steps. Therefore in adding your new option, be sure to keep it in the same relative place in the list in all the separate places (it doesn't necessarily have to be in the end), and near conceptually similar options.

`main.h`
: The `objectcols` and `clumpcols` enumerated variables (`enum`) define the raw/internal calculation columns. If your new column requires new raw calculations, add a row to the respective list. If your calculation requires any other settings paramters, you should add a variable to the `mkcatalogparams` structure.

`ui.h` The `option_keys_enum` associates a unique value for each option to MakeProfiles. The options that have a short option version, the single character short comment is used for the value. Those that don't have a short option version, get a large integer automatically. You should add a variable here to identify your desired column.

`args.h` This file specifies all the parameters for the GNU C library, Argp structure that is in charge of reading the user's options. To define your new column, just copy an existing set of parameters and change the first, second and 5th values (the only ones that differ between all the columns), you should use the macro you defined in `ui.h` here.

`ui.c` If your column includes any particular settings (you added a variable to the `mkcatalogparams` structure in `main.h`), you should do the sanity checks and preparations for it here. Otherwise, you can ignore this file.

`columns.c`
This file will contain the main defintion and high-level calculation of your new column through the `columns_define_alloc` and `columns_fill` functions. In the first, you specify the basic information about the column: its name, units, comments, type (see Section 4.4 [Numeric data types], page 64) and how it should be printed if the output is a text file. You should also specify the raw/internal columns that are necessary for this column here as the many existing examples show. Through the types for objects and rows, you can specify if this column is only for clumps, objects or both.

The second main function (`columns_fill`) writes the final value into the appropriate column for each object and clump. As you can see in the many existing examples, you can define your processing on the raw/internal calculations here and save them in the output.

`mkcatalog.c`
As described before, this file contains the two main MakeCatalog work-horses: `mkcatalog_first_pass` and `mkcatalog_second_pass`, their names are descriptive enough and their internals are also clear and heavily commented.

7.3.5 Invoking MakeCatalog

MakeCatalog will make a catalog from an input image and at least on labeled image. The executable name is `astmkcatalog` with the following general template

```
$ astmkcatalog [OPTION ...] InputImage.fits
```

One line examples:

```
## Create catalog with RA, Dec, Magnitude and Magnitude error,
## 'input.fits' is NoiseChisel's output:
$ astmkcatalog --ra --dec --magnitude --magnitudeerr input.fits

## Same catalog as above (using short options):
$ asmkcatalog -rdmG input.fits

## Write the catalog to a FITS table:
```

```
$ astmkcatalog -mpQ --output=cat.fits input_labeled.fits

## Read the columns to create from 'columns.conf':
$ astmkcatalog --config=columns.conf input_labeled.fits

## Use different images for the objects and clumps inputs:
$ astmkcatalog --objectsfile=K_labeled.fits --objectshdu=1    \
               --clumpsfile=K_labeled.fits --clumpshdu=2 i_band.fits
```

If MakeCatalog is to do processing, an input image should be provided with the recognized extensions as input data, see Section 4.1.1.1 [Arguments], page 49. The options described in this section are those that are only particular to MakeProfiles. For operations that MakeProfiles shares with other programs (mainly involving input/output or general processing steps), see Section 4.1.2 [Common options], page 52. Also see Chapter 4 [Common program behavior], page 48, for some general characteristics of all Gnuastro programs including MakeCatalog.

MakeCatalog needs 4 (or 5) images as input. These images can be separate extensions in one file (NoiseChisel's default output), or each can have its own file and its own extension. See Section 7.2.2.4 [NoiseChisel output], page 175, for the list. The clump labels image is not mandatory (when no clump catalog is required, for example in aperture photometry). When inspecting the object labels image, MakeProfiles will look for a `WCLUMPS` (short for with-clumps) header keyword. If that keyword is present and has a value of `yes`, `1`, or `y` (case insensitive) then a clump image must also be provided and a clump catalog will be made. When `WCLUMPS` isn't present or has any other value, only an object catalog will be created and all clump related options/columns will be ignored.

For example, if you only need an object catalog from NoiseChisel's output, you can use Gnuastro's Fits program (see Section 5.1 [Fits], page 80) to modify or remove the `WCLUMPS` keyword in the objects HDU, then run MakeCatalog on it. Another example can be aperture photometry: let's assume you have made your labeled image (defining the apertures) with MakeProfiles. Clumps are not defined in this context, so besides the input and labeled image, you only need NoiseChisel's Sky and Sky standard deviation images (run NoiseChisel with the `--onlydetection` option). Since MakeProfile's output doesn't contain the `WCLUMPS` keyword, you just have to specify your labeled image with the `--objectsfile` option and also set its HDU. Note that labeled images have to be an integer type. Therefore, if you are using MakeProfiles to define the apertures/labels, you can use its `--type=int32` for example, see Section 4.1.2.1 [Input/Output options], page 52, and Section 4.4 [Numeric data types], page 64.

When a clump catalog is also desired, two catalogs will be made: one for the objects (suffixed with `_o.txt` or `_o.fits`) and another for the clumps (suffixed with `_c.txt` or `_c.fits`). Therefore if any value is given to the `--output` option, MakeCatalogs will replace these two suffixes with any existing suffix in the given value. If no output value is given, MakeCatalog will use the input name, see Section 4.8 [Automatic output], page 77. The format of the output table is specified with the `--tableformat` option, see Section 4.1.2.1 [Input/Output options], page 52.

When MakeCatalog is run on multiple threads, the clumps catalog rows will not be sorted by object since each object is processed independently by one thread and threaded applications are asynchronous. The clumps in each object will be sorted based on their

labels, but you will find lower-index objects come after higher-index ones (especially if they have more clumps and thus take more time). If the order is very important for you, you can run the following command to sort the rows by object ID (and clump ID with each object):

```
$ awk '!/^#/' out_c.txt | sort -g -k1,1 -k2,2
```

7.3.5.1 MakeCatalog input files

MakeCatalog needs multiple images as input: a values image, one (or two) labeled images and Sky and Sky standandard deviation images. The options described in this section allow you to identify them. If you use the default output of NoiseChisel (see Section 7.2.2.4 [NoiseChisel output], page 175) you don't have to worry about any of these options and just give NoiseChisel's output file to MakeCatalog as described in Section 7.3.5 [Invoking MakeCatalog], page 185.

`-O STR`
`--objectsfile=STR`
: The file name of the object labels image, if the image is in another extension of the input file, calling this option is not mandatory, just specify the extension/HDU with the `--objectshdu` option.

`--objectshdu=STR`
: The HDU/extension of the object labels image. Only pixels with values above zero will be considered. The objects label image has to be an integer data type (see Section 4.4 [Numeric data types], page 64) and only pixels with a value larger than zero will be used. If this extension contains teh `WCLUMPS` keyword with a value of `yes`, `1`, or `y` (not case sensitive), then MakeCatalog will also build a clumps catalog, see Section 7.3.5 [Invoking MakeCatalog], page 185.

`-C STR`
`--clumpsfile=STR`
: Similar to `--objlabs` but for the labels of the clumps. This is only necessary if the image containing clump labels is not in the input file and the objects image has a `WCLUMPS` keyword, see `--objectshdu`.

`--clumpshdu=STR`
: The HDU/extension of the object labels image. Only pixels with values above zero will be considered. The objects label image has to be an integer data type (see Section 4.4 [Numeric data types], page 64) and only pixels with a value larger than zero will be used.

`-s STR`
`--skyfile=STR`
: File name of an image keeping the Sky value for each pixel.

`--skyhdu=STR`
: The HDU of the Sky value image.

`-t STR`
`--stdfile=STR`
: File name of image keeping the Sky value standard deviation for each pixel.

`--stdhdu=STR`
: The HDU of the Sky value standard deviation image.

7.3.5.2 MakeCatalog general settings

Some of the columns require particular settings (for example the zero point magnitdue for measuring magnitudes), the options in this section can be used for such configurations.

`-z FLT`
`--zeropoint=FLT`
: The zero point magnitude for the input image, see Section 8.1.3 [Flux Brightness and magnitude], page 201.

`-E`
`--skysubtracted`
: If the image has already been sky subtracted by another program, then you need to notify MakeCatalog through this option. Note that this is only relevant when the Signal to noise ratio is to be calculated.

`-T FLT`
`--threshold=FLT`
: For all the columns, only consider pixels that are above a given relative threshold. Symbolizing the value of this option as T, the Sky for a pixel at (i, j) with μ_{ij} and its Standard deviation with σ_{ij}, that pixel will only be used if its value (B_{ij}) satisfies this condition: $B_{ij} > \mu_{ij} + T\sigma_{ij}$. The only calculations that will not be affected are is the average river values (`--riverave`), since they are used as a reference. A commented row will be added in the header of the output catalog that will print the given value, since this is a very important issue, it starts with `**IMPORTANT**`.

  NoiseChisel will detect very diffuse signal which is useful in most cases where the aggregate properties of the detections are desired, since there is signal there (with the desired certainty). However, in some cases, only the properties of the peaks of the objects/clumps are desired, for example in attempting to separate stars from galaxies, the peaks are the major target and the diffuse regions only act to complicate the separation. With this option, MakeCatalog will simply ignore any pixel below the relative threshold.

  This option is not mandatory, so if it isn't given (after reading the command-line and all configuration files, see Section 4.2 [Configuration files], page 59), MakeCatalog will still operate. However, if it has a value in any lower-level configuration file and you want to ignore that value for this particular run or in a higher-level configuration file, then set it to NaN, for example `--threshold=nan`. Gnuastro uses the C library's `strtod` function to read floats, which is not case-sensative in reading NaN values. But to be consistent, it is good practice to only use `nan`.

`--nsigmag=FLT`
: The median standard deviation (from the standard deviation image) will be multiplied by the value to this option and its magnitude will be reported in the comments of the output catalog. This value is a per-pixel value, not per object/clump and is not found over an area or aperture, like the common 5σ values that are commonly reported as a measure of depth or the upper-limit measurements (see Section 7.3.2 [Quantifying measurement limits], page 177).

7.3.5.3 Upper-limit magnitude settings

The upper limit magnitude was discussed in Section 7.3.2 [Quantifying measurement limits], page 177. Unlike other measured values/columns in MakeCatalog, the upper limit magnitude needs several defined parameters which are discussed here. All the upper limit magnitude specific options start with `up` for upper-limit, except for `--envseed` that is also present in other programs and is general for any job requiring random number generation (see Section 8.2.1.4 [Generating random numbers], page 212).

One very important consideration in Gnuastro is reproducibility. Therefore, the values to all of these parameters along with others (like the random number generator type and seed) are also reported in the comments of the final catalog when the upper limit magnitude column is desired. The random seed that is used to define the random positionings for each object or clump is unique and set based on the given seed, the total number of objects and clumps and also the labels of the clumps and objects. So with identical inputs, an identical upper-limit magnitude will be found. But even if the ordering of the object/clump labels differs (and the seed is the same) the result will not be the same.

MakeCatalog will randomly place the object/clump footprint over the image and when the footprint doesn't fall on any object or masked region (see `--upmaskfile`) it will be used until the desired number (`--upnum`) of samples are found to estimate the distribution's standard deviation (see Section 7.3.2 [Quantifying measurement limits], page 177). Otherwise it will be ignored and another random position will be generated. But when the profile is very large or the image is significantly covered by detections, it might not be possible to find the desired number of samplings. MakeProfiles will continue searching until 50 times the value given to `--upnum`. If `--upnum` good samples cannot be found until this limit, it will set the upper-limit magnitude for that object to NaN (blank).

`--upmaskfile=STR`
: File name of mask image to use for upper-limit calculation. In some cases (especially when doing matched photometry), the object labels specified in the main input and mask image might not be adequate. In other words they do not necessarily have to cover *all* detected objects: the user might have selected only a few of the objects in their labeled image. This option can be used to ignore regions in the image in these situations when estimating the upper-limit magnitude. All the non-zero pixels of the image specified by this option (in the `--upmaskhdu` extension) will be ignored in the upper-limit magnitude measurements.

  For example, when you are using labels from another image, you can give NoiseChisel's objects image output for this image as the value to this option. In this way, you can be sure that regions with data do not harm your distribution. See Section 7.3.2 [Quantifying measurement limits], page 177, for more on the upper limit magnitude.

`--upmaskhdu=STR`
: The extension in the file specified by `--upmask`.

`--upnum=INT`
: The number of random samples to take for all the objects. A larger value to this option will give a more accurate result (assymptotically), but it will also slow down the process. When a randomly positioned sample overlaps with a

detected/masked pixel it is not counted and another random position is found until the object completely lies over an undetected region. So you can be sure that for each object, this many samples over undetected objects are made. See the upper limit magnitude discussion in Section 7.3.2 [Quantifying measurement limits], page 177, for more.

`--envseed`
: Read the random number generator type and seed value from the environment (see Section 8.2.1.4 [Generating random numbers], page 212). Random numbers are used in calculating the random positions of different samples of each object.

`--upsigmaclip=FLT,FLT`
: The raw distribution of random values will not be used to find the upper-limit magnitude, it will first be σ-clipped (see Section 7.1.2 [Sigma clipping], page 149) to avoid outliers in the distribution (mainly the faint undetected wings of bright/large objects in the image). This option takes two values: the first is the multiple of σ, and the second is theh termination criteria. If the latter is larger than 1, it is read as an integer number and will be the number of times to clip. If it is smaller than 1, it is interpretted as the tollerance level to stop clipping. See Section 7.1.2 [Sigma clipping], page 149, for a complete explanation.

`--upnsigma=FLT`
: The multiple of the final (σ-clipped) standard deviation (or σ) used to measure the upper-limit brightness or magnitude.

7.3.5.4 MakeCatalog output columns

The final group of options particular to MakeCatalog are those that specify which columns should be written into the final output table. For each column there is an option, if it has been called on the command line or in any of the configuration files, it will included as a column in the output catalog in the same order (see Section 4.2.2 [Configuration file precedence], page 60). Some of the columns apply to both objects and clumps and some are particular to only one of them. The latter cases are explicitly marked with [Objects] or [Clumps] to specify the catalog they will be placed in.

`--i`
`--ids`
: This is a unique option it can add multiple columns to the final catalog(s). Calling this option will put the object IDs (`--objid`) in the objects catalog and host-object-ID (`--hostobjid`) and ID-in-host-object (`--idinhostobj`) into the clumps catalog. Hence if only object catalogs are required, it has the same effect as `--objid`.

`--objid`
: [Objects] ID of this object.

`-j`
`--hostobjid`
: [Clumps] The ID of the object which hosts this clump.

`--idinhostobj`
: [Clumps] The ID of this clump in its host object.

`-C`
`--numclumps`
[Objects] The number of clumps in this object.

`-a`
`--area` The raw area (number of pixels) in any clump or object independed of what pixel it lies over (if it is NaN/blank or unused for example).

`--clumpsarea`
[Objects] The total area of all the clumps in this object.

`--weightarea`
The area (number of pixels) used in the flux weighted position calculations.

`-x`
`--x` The flux weighted center of all objects and clumps along the first FITS axis (horizontal when viewed in SAO ds9), see $\overline{x}$ in Section 7.3.3 [Measuring elliptical parameters], page 181. The weight has to have a positive value (pixel value larger than the Sky value) to be meaningful! Specially when doing matched photometry, this might not happen: no pixel value might be above the Sky value. For such detections, the geometric center will be reported in this column (see `--geox`). You can use `--weightarea` to see which was used.

`-y`
`--y` The flux weighted center of all objects and clumps along the second FITS axis (vertical when viewed in SAO ds9). See `--x`.

`--geox` The geometric center of all objects and clumps along the first FITS axis axis. The geometric center is the average pixel positions irrespective of their pixel values.

`--geoy` The geometric center of all objects and clumps along the second FITS axis axis, see `--geox`.

`--clumpsx`
[Objects] The flux weighted center of all the clumps in this object along the first FITS axis. See `--x`.

`--clumpsy`
[Objects] The flux weighted center of all the clumps in this object along the second FITS axis. See `--x`.

`--clumpsgeox`
[Objects] The geometric center of all the clumps in this object along the first FITS axis. See `--geox`.

`--clumpsgeoy`
[Objects] The geometric center of all the clumps in this object along the second FITS axis. See `--geox`.

`-r`
`--ra` Flux weighted right ascension of all objects or clumps, see `--x`.

`-d`
`--dec` Flux weighted declination of all objects or clumps, see `--x`.

`--geora` Geometric center right ascension of all objects or clumps, see `--geox`.

`--geodec` Geometric center declination of all objects or clumps, see `--geox`.

`--clumpsra`
: [Objects] Flux weighted right ascension of all clumps in this object, see `--x`.

`--clumpsdec`
: [Objects] Flux weighted declination of all clumps in this object, see `--x`.

`--clumpsgeora`
: [Objects] Geometric center right ascension of all clumps in this object, see `--geox`.

`--clumpsgeodec`
: [Objects] Geometric center declination of all clumps in this object, see `--geox`.

`-b`
`--brightness`
: The brightness (sum of all pixel values), see Section 8.1.3 [Flux Brightness and magnitude], page 201. For clumps, the ambient brightness (flux of river pixels around the clump multiplied by the area of the clump) is removed, see `--riverflux`. So the sum of clump brightnesses in the clump catalog will be smaller than the total clump brightness in the `--clumpbrightness` column of the objects catalog.

  If no usable pixels (blank or below the threshold) are present over the clump or object, the stored value will be NaN (note that zero is meaningful).

`--clumpbrightness`
: [Objects] The total brightness of the clumps within an object. This is simply the sum of the pixels associated with clumps in the object. If no usable pixels (blank or below the threshold) are present over the clump or object, the stored value will be NaN, because zero (note that zero is meaningful).

`--noriverbrightness`
: [Clumps] The Sky (not river) subtracted clump brightness. By definition, for the clumps, the average brightness of the rivers surrounding it are subtracted from it for a first order accounting for contamination by neighbors. In cases where you will be calculating the flux brightness difference later (one example below) the contamination will be (mostly) removed at that stage, which is why this column was added.

  One example might be this: you want to know the change in the clump flux as a function of threshold (see `--threshold`). So you will make two catalogs (each having this column but with different thresholds) and then subtract the lower threshold catalog (higher brightness) from the higher threshold catalog (lower brightness). The effect is most visible when the rivers have a high average signal-to-noise ratio. The removed contribution from the pixels below the threshold will be less than the river pixels. Therefore the river-subtracted brightness (`--brightness`) for the thresholded catalog for such clumps will be larger than the brightness with no threshold!

  If no usable pixels (blank or below the possibly given threshold) are present over the clump or object, the stored value will be NaN (note that zero is meaningful).

`-m`
`--magnitude`
The magnitude of clumps or objects, see `--brightness`.

`-e`
`--magnitudeerr`
The magnitude error of clumps or objects. The magnitude error is calculated from the signal-to-noise ratio (see `--sn` and Section 7.3.2 [Quantifying measurement limits], page 177). Note that until now this error assumes un-correlated pixel values and also does not include the error in estimating the aperture (or error in generating the labeled image).

For now these factors have to be found by other means. Task 14124 (`https://savannah.gnu.org/task/index.php?14124`) has been defined for work on adding these sources of error too.

`--clumpsmagnitude`
[Objects] The magnitude of all clumps in this object, see `--clumpbrightness`.

`--upperlimit`
The upper limit value (in units of the input image) for this object or clump. See Section 7.3.2 [Quantifying measurement limits], page 177, and Section 7.3.5.3 [Upper-limit magnitude settings], page 189, for a complete explanation. This is very important for the fainter and smaller objects in the image where the measured magnitudes are not reliable.

`--upperlimitmag`
The upper limit magnitude for this object or clump. See Section 7.3.2 [Quantifying measurement limits], page 177, and Section 7.3.5.3 [Upper-limit magnitude settings], page 189, for a complete explanation. This is very important for the fainter and smaller objects in the image where the measured magnitudes are not reliable.

`--riverave`
[Clumps] The average brightness of the river pixels around this clump. River pixels were defined in Akhlaghi and Ichikawa 2015. In short they are the pixels immediately outside of the clumps. This value is used internally to find the brightness (or magnitude) and signal to noise ratio of the clumps. It can generally also be used as a scale to gauge the base (ambient) flux surrounding the clump. In case there was no river pixels, then this column will have the value of the Sky under the clump. So note that this value is *not* sky subtracted.

`--rivernum`
[Clumps] The number of river pixels around this clump, see `--riverflux`.

`-n`
`--sn` The Signal to noise ratio (S/N) of all clumps or objects. See Akhlaghi and Ichikawa (2015) for the exact equations used.

`--sky` The sky flux (per pixel) value under this object or clump. This is actually the mean value of all the pixels in the sky image that lie on the same position as the object or clump.

`--std` The sky value standard deviation (per pixel) for this clump or object. Like `--sky`, this is the average of the values in the input sky standard deviation image pixels that lie over this object.

`-A`
`--semimajor`
The pixel-value weighted semi-major axis of the profile (assuming it is an ellipse) in units of pixels. See Section 7.3.3 [Measuring elliptical parameters], page 181.

`-B`
`--semiminor`
The pixel-value weighted semi-minor axis of the profile (assuming it is an ellipse) in units of pixels. See Section 7.3.3 [Measuring elliptical parameters], page 181.

`-p`
`--positionangle`
The pixel-value weighted angle of the semi-major axis with the first FITS axis in degrees. See Section 7.3.3 [Measuring elliptical parameters], page 181.

`--geosemimajor`
The geometric (ignoring pixel values) semi-major axis of the profile, assuming it is an ellipse.

`--geosemiminor`
The geometric (ignoring pixel values) semi-minor axis of the profile, assuming it is an ellipse.

`--geopositionangle`
The geometric (ignoring pixel values) angle of the semi-major axis with the first FITS axis in degrees.

8 Modeling and fitting

In order to fully understand observations after initial analysis on the image, it is very important to compare them with the existing models to be able to further understand both the models and the data. The tools in this chapter create model galaxies and will provide 2D fittings to be able to understand the detections.

8.1 MakeProfiles

MakeProfiles will create mock astronomical profiles from a catalog, either individually or together in one output image. In data analysis, making a mock image can act like a calibration tool, through which you can test how successfully your detection technique is able to detect a known set of objects. There are commonly two aspects to detecting: the detection of the fainter parts of bright objects (which in the case of galaxies fade into the noise very slowly) or the complete detection of an over-all faint object. Making mock galaxies is the most accurate (and idealistic) way these two aspects of a detection algorithm can be tested. You also need mock profiles in fitting known functional profiles with observations.

MakeProfiles was initially built for extra galactic studies, so currently the only astronomical objects it can produce are stars and galaxies. We welcome the simulation of any other astronomical object. The general outline of the steps that MakeProfiles takes are the following:

1. Build the full profile out to its truncation radius in a possibly over-sampled array.
2. Multiply all the elements by a fixed constant so its total magnitude equals the desired total magnitude.
3. If `--individual` is called, save the array for each profile to a FITS file.
4. If `--nomerged` is not called, add the overlapping pixels of all the created profiles to the output image and abort.

Using input values, MakeProfiles adds the World Coordinate System (WCS) headers of the FITS standard to all its outputs (except PSF images!). For a simple test on a set of mock galaxies in one image, there is no need for the third step or the WCS information.

However in complicated simulations like weak lensing simulations, where each galaxy undergoes various types of individual transformations based on their position, those transformations can be applied to the different individual images with other programs. After all the transformations are applied, using the WCS information in each individual profile image, they can be merged into one output image for convolution and adding noise.

8.1.1 Modeling basics

In the subsections below, first a review of some very basic information and concepts behind modeling a real astronomical image is given. You can skip this subsection if you are already sufficiently familiar with these concepts.

8.1.1.1 Defining an ellipse

The PSF, see Section 8.1.1.2 [Point Spread Function], page 196, and galaxy radial profiles are generally defined on an ellipse so in this section first defining an ellipse on a pixelated 2D surface is discussed. Labeling the major axis of an ellipse a, and its minor axis with b,

the axis ratio is defined as: $q \equiv b/a$. The major axis of an ellipse can be aligned in any direction, therefore define the angle of the major axis to the horizontal axis of the image is defined to be the position angle of the ellipse and in this book, we show it with θ.

Our aim is to put a radial profile of any functional form $f(r)$ over an ellipse. Let's define the radial distance r_{el} as the distance on the major axis to the center of the ellipse which is located at x_c and y_c. We want to find the elliptical distance of a point located at (i, j), in the image coordinate system, from the center of the ellipse. First the coordinate system is rotated by θ to get the new rotated coordinates of that point (i_r, j_r):

$$i_r(i,j) = (i_c - i)\cos(\theta) + (j_c - j)\sin(\theta)$$

$$j_r(i,j) = (j_c - j)\cos(\theta) - (i_c - i)\sin(\theta)$$

The elliptical distance of a point located at (i, j) can now be defined as: $r_{el}^2 = \sqrt{i_r^2 + j_r^2/q^2}$. To place the radial profiles explained below over an ellipse, $f(r_{el}(i, j))$ is calculated based on the functional radial profile desired.

The way MakeProfiles builds the profile is that the nearest pixel in the image to the given profile center is found and the profile value is calculated for it, see Section 8.1.1.5 [Sampling from a function], page 199. The next pixel which the profile value is calculated on is the next nearest neighbor of the initial pixel to the profile center (as defined by r_{el}). This is done fairly efficiently using a breadth first parsing strategy[1] which is implemented through an ordered linked list.

Using this approach, we only go over one layer of pixels on the circumference of the profile to build the profile. Not one more extra pixel has to be checked. Another consequence of this strategy is that extending MakeProfiles to three dimensions becomes very simple: only the neighbors of each pixel have to be changed. Everything else after that (when the pixel index and its radial profile have entered the linked list) is the same, no matter the number of dimensions we are dealing with.

8.1.1.2 Point Spread Function

Assume we have a 'point' source, or a source that is far smaller than the maximum resolution (a pixel). When we take an image of it, it will 'spread' over an area. To quantify that spread, we can define a 'function'. This is how the point spread function or the PSF of an image is defined. This 'spread' can have various causes, for example in ground based astronomy, due to the atmosphere. In practice we can never surpass the 'spread' due to the diffraction of the lens aperture. Various other effects can also be quantified through a PSF. For example, the simple fact that we are sampling in a discrete space, namely the pixels, also produces a very small 'spread' in the image.

Convolution is the mathematical process by which we can apply a 'spread' to an image, or in other words blur the image, see Section 6.3.1.1 [Convolution process], page 118. The Brightness of an object should remain unchanged after convolution, see Section 8.1.3 [Flux Brightness and magnitude], page 201. Therefore, it is important that the sum of all the pixels of the PSF be unity. The PSF image also has to have an odd number of pixels on its

---
[1] `http://en.wikipedia.org/wiki/Breadth-first_search`

sides so one pixel can be defined as the center. In MakeProfiles, the PSF can be set by the two methods explained below.

Parametric functions

A known mathematical function is used to make the PSF. In this case, only the parameters to define the functions are necessary and MakeProfiles will make a PSF based on the given parameters for each function. In both cases, the center of the profile has to be exactly in the middle of the central pixel of the PSF (which is automatically done by MakeProfiles). When talking about the PSF, usually, the full width at half maximum or FWHM is used as a scale of the width of the PSF.

Gaussian In the older papers, and to a lesser extent even today, some researchers use the 2D Gaussian function to approximate the PSF of ground based images. In its most general form, a Gaussian function can be written as:

$$f(r) = a \exp\left(\frac{-(x-\mu)^2}{2\sigma^2}\right) + d$$

Since the center of the profile is pre-defined, μ and d are constrained. a can also be found because the function has to be normalized. So the only important parameter for MakeProfiles is the σ. In the Gaussian function we have this relation between the FWHM and σ:

$$\mathrm{FWHM_g} = 2\sqrt{2\ln 2}\sigma \approx 2.35482\sigma$$

Moffat The Gaussian profile is much sharper than the images taken from stars on photographic plates or CCDs. Therefore in 1969, Moffat proposed this functional form for the image of stars:

$$f(r) = a\left[1 + \left(\frac{r}{\alpha}\right)^2\right]^{-\beta}$$

Again, a is constrained by the normalization, therefore two parameters define the shape of the Moffat function: α and β. The radial parameter is α which is related to the FWHM by

$$\mathrm{FWHM_m} = 2\alpha\sqrt{2^{1/\beta} - 1}$$

Comparing with the PSF predicted from atmospheric turbulence theory with a Moffat function, Trujillo et al.[2] claim that β should

[2] Trujillo, I., J. A. L. Aguerri, J. Cepa, and C. M. Gutierrez (2001). “The effects of seeing on Sérsic profiles - II. The Moffat PSF”. In: MNRAS 328, pp. 977—985.

be 4.765. They also show how the Moffat PSF contains the Gaussian PSF as a limiting case when $\beta \to \infty$.

An input FITS image
: An input image file can also be specified to be used as a PSF. If the sum of its pixels are not equal to 1, the pixels will be multiplied by a fraction so the sum does become 1.

While the Gaussian is only dependent on the FWHM, the Moffat function is also dependent on β. Comparing these two functions with a fixed FWHM gives the following results:

- Within the FWHM, the functions don't have significant differences.
- For a fixed FWHM, as β increases, the Moffat function becomes sharper.
- The Gaussian function is much sharper than the Moffat functions, even when β is large.

8.1.1.3 Stars

In MakeProfiles, stars are generally considered to be a point source. This is usually the case for extra galactic studies, were nearby stars are also in the field. Since a star is only a point source, we assume that it only fills one pixel prior to convolution. In fact, exactly for this reason, in astronomical images the light profiles of stars are one of the best methods to understand the shape of the PSF and a very large fraction of scientific research is preformed by assuming the shapes of stars to be the PSF of the image.

8.1.1.4 Galaxies

Today, most practitioners agree that the flux of galaxies can be modeled with one or a few generalized de Vaucouleur's (or Sérsic) profiles.

$$I(r) = I_e \exp\left(-b_n\left[\left(\frac{r}{r_e}\right)^{1/n} - 1\right]\right)$$

Gérard de Vaucouleurs (1918-1995) was first to show in 1948 that this function best fits the galaxy light profiles, with the only difference that he held n fixed to a value of 4. 20 years later in 1968, J. L. Sérsic showed that n can have a variety of values and does not necessarily need to be 4. This profile depends on the effective radius (r_e) which is defined as the radius which contains half of the profile brightness (see Section 8.1.4 [Profile magnitude], page 201). I_e is the flux at the effective radius. The Sérsic index n is used to define the concentration of the profile within r_e and b_n is a constant dependent on n. MacArthur et al.[3] show that for $n > 0.35$, b_n can be accurately approximated using this equation:

$$b_n = 2n - \frac{1}{3} + \frac{4}{405n} + \frac{46}{25515n^2} + \frac{131}{1148175n^3} - \frac{2194697}{30690717750n^4}$$

[3] MacArthur, L. A., S. Courteau, and J. A. Holtzman (2003). "Structure of Disk-dominated Galaxies. I. Bulge/Disk Parameters, Simulations, and Secular Evolution". In: ApJ 582, pp. 689—722.

8.1.1.5 Sampling from a function

A pixel is the ultimate level of accuracy to gather data, we can't get any more accurate in one image, this is known as sampling in signal processing. However, the mathematical profiles which describe our models have infinite accuracy. Over a large fraction of the area of astrophysically interesting profiles (for example galaxies or PSFs), the variation of the profile over the area of one pixel is not too significant. In such cases, the elliptical radius (r_{el} of the center of the pixel can be assigned as the final value of the pixel, see Section 8.1.1.1 [Defining an ellipse], page 195).

As you approach their center, some galaxies become very sharp (their value significantly changes over one pixel's area). This sharpness increases with smaller effective radius and larger Sérsic values. Thus rendering the central value extremely inaccurate. The first method that comes to mind for solving this problem is integration. The functional form of the profile can be integrated over the pixel area in a 2D integration process. However, unfortunately numerical integration techniques also have their limitations and when such sharp profiles are needed they can become extremely inaccurate.

The most accurate method of sampling a continuous profile on a discrete space is by choosing a large number of random points within the boundaries of the pixel and taking their average value (or Monte Carlo integration). This is also, generally speaking, what happens in practice with the photons on the pixel. The number of random points can be set with `--numrandom`.

Unfortunately, repeating this Monte Carlo process would be extremely time and CPU consuming if it is to be applied to every pixel. In order to not loose too much accuracy, in MakeProfiles, the profile is built using both methods explained below. The building of the profile begins from its central pixel and continues (radially) outwards. Monte Carlo integration is first applied (which yields F_r), then the central pixel value (F_c) is calculated on the same pixel. If the fractional difference ($|F_r - F_c|/F_r$) is lower than a given tolerance level (specified with `--tolerance`) MakeProfiles will stop using Monte Carlo integration and only use the central pixel value.

The ordering of the pixels in this inside-out construction is based on $r = \sqrt{(i_c - i)^2 + (j_c - j)^2}$, not r_{el}, see Section 8.1.1.1 [Defining an ellipse], page 195. When the axis ratios are large (near one) this is fine. But when they are small and the object is highly elliptical, it might seem more reasonable to follow r_{el} not r. The problem is that the gradient is stronger in pixels with smaller r (and larger r_{el}) than those with smaller r_{el}. In other words, the gradient is strongest along the minor axis. So if the next pixel is chosen based on r_{el}, the tolerance level will be reached sooner and lots of pixels with large fractional differences will be missed.

Monte Carlo integration uses a random number of points. Thus, every time you run it, by default, you will get a different distribution of points to sample within the pixel. In the case of large profiles, this will result in a slight difference of the pixels which use Monte Carlo integration each time MakeProfiles is run. To have a deterministic result, you have to fix the random number generator properties which is used to build the random distribution. This can be done by setting the `GSL_RNG_TYPE` and `GSL_RNG_SEED` environment variables and calling MakeProfiles with the `--envseed` option. To learn more about the process of generating random numbers, see Section 8.2.1.4 [Generating random numbers], page 212.

The seed values are fixed for every profile: with `--envseed`, all the profiles have the same seed and without it, each will get a different seed using the system clock (which is accurate to within one microsecond). The same seed will be used to generate a random number for all the sub-pixel positions of all the profiles. So in the former, the sub-pixel points checked for all the pixels undergoing Monte carlo integration in all profiles will be identical. In other words, the sub-pixel points in the first (closest to the center) pixel of all the profiles will be identical with each other. All the second pixels studied for all the profiles will also receive an identical (different from the first pixel) set of sub-pixel points and so on. As long as the number of random points used is large enough or the profiles are not identical, this should not cause any systematic bias.

8.1.1.6 Oversampling

The steps explained in Section 8.1.1.5 [Sampling from a function], page 199, do give an accurate representation of a profile prior to convolution. However, in an actual observation, the image is first convolved with or blurred by the atmospheric and instrument PSF in a continuous space and then it is sampled on the discrete pixels of the camera.

In order to more accurately simulate this process, the unconvolved image and the PSF are created on a finer pixel grid. In other words, the output image is a certain odd-integer multiple of the desired size, we can call this 'oversampling'. The user can specify this multiple as a command-line option. The reason this has to be an odd number is that the PSF has to be centered on the center of its image. An image with an even number of pixels on each side does not have a central pixel.

The image can then be convolved with the PSF (which should also be oversampled on the same scale). Finally, image can be sub-sampled to get to the initial desired pixel size of the output image. After this, mock noise can be added as explained in the next section. This is because unlike the PSF, the noise occurs in each output pixel, not on a continuous space like all the prior steps.

8.1.2 If convolving afterwards

In case you want to convolve the image later with a given point spread function, make sure to use a larger image size. After convolution, the profiles become larger and a profile that is normally completely outside of the image might fall within it.

On one axis, if you want your final (convolved) image to be m pixels and your PSF is $2n+1$ pixels wide, then when calling MakeProfiles, set the axis size to $m+2n$, not m. You also have to shift all the pixel positions of the profile centers on the that axis by n pixels to the positive.

After convolution, you can crop the outer n pixels with the section crop box specification of Crop: `--section=n:*-n,n:*-n` assuming your PSF is a square, see Section 6.1.2 [Crop section syntax], page 100. This will also remove all discrete Fourier transform artifacts (blurred sides) from the final image. To facilitate this shift, MakeProfiles has the options `--xshift`, `--yshift` and `--prepforconv`, see Section 8.1.5 [Invoking MakeProfiles], page 202.

8.1.3 Flux Brightness and magnitude

Astronomical data pixels are usually in units of counts[4] or electrons or either one divided by seconds. To convert from the counts to electrons, you will need to know the instrument gain. In any case, they can be directly converted to energy or energy/time using the basic hardware (telescope, camera and filter) information. We will continue the discussion assuming the pixels are in units of energy/time.

The *brightness* of an object is defined as its total detected energy per time. This is simply the sum of the pixels that are associated with that detection by our detection tool for example Section 7.2 [NoiseChisel], page 163[5]. The *flux* of an object is in units of energy/time/area and for a detected object, it is defined as its brightness divided by the area used to collect the light from the source or the telescope aperture (for example in cm^2)[6]. Knowing the flux (f) and distance to the object (r), we can calculate its *luminosity*: $L = 4\pi r^2 f$. Therefore, flux and luminosity are intrinsic properties of the object, while brightness depends on our detecting tools (hardware and software). Here we will not be discussing luminosity, but brightness. However, since luminosity is the astrophysically interesting quantity, we also defined it here to avoid possible confusion between these two terms because they both have the same units.

Images of astronomical objects span over a very large range of brightness. With the Sun (as the brightest object) being roughly $2.5^{60} = 10^{24}$ times brighter than the faintest galaxies we can currently detect. Therefore discussing brightness will be very hard, and astronomers have chosen to use a logarithmic scale to talk about the brightness of astronomical objects. But the logarithm can only be usable with a unit-less and always positive value. Fortunately brightness is always positive and to remove the units we divide the brightness of the object (B) by a reference brightness (B_r). We then define the resulting logarithmic scale as *magnitude* through the following relation[7]

$$m - m_r = -2.5\log_{10}\left(\frac{B}{B_r}\right)$$

m is defined as the magnitude of the object and m_r is the pre-defined magnitude of the reference brightness. One particularly easy condition is when $B_r = 1$. This will allow us to summarize all the hardware specific parameters discussed above into one number as the reference magnitude which is commonly known as the Zero-point[8] magnitude.

8.1.4 Profile magnitude

To find the profile brightness or its magnitude, (see Section 8.1.3 [Flux Brightness and magnitude], page 201), it is customary to use the 2D integration of the flux to infinity.

[4] Counts are also known as analog to digital units (ADU).

[5] If further processing is done, for example the Kron or Petrosian radii are calculated, then the detected area is not sufficient and the total area that was within the respective radius must be used.

[6] For a full object that spans over several pixels, the telescope area should be used to find the flux. However, sometimes, only the brightness per pixel is desired. In such cases this book also *loosely* uses the term flux. This is only approximately accurate however, since while all the pixels have a fixed area, the pixel size can vary with camera on the telescope.

[7] The -2.5 factor in the definition of magnitudes is a legacy of the our ancient colleagues and in particular Hipparchus of Nicaea (190-120 BC).

[8] When $B = Br = 1$, the right side of the magnitude definition will be zero. Hence the name, “zero-point”.

However, in MakeProfiles we do not follow this idealistic approach and apply a more realistic method to find the total brightness or magnitude: the sum of all the pixels belonging to a profile within its predefined truncation radius. Note that if the truncation radius is not large enough, this can be significantly different from the total integrated light to infinity.

An integration to infinity is not a realistic condition because no galaxy extends indefinitely (important for high Sérsic index profiles), pixelation can also cause a significant difference between the actual total pixel sum value of the profile and that of integration to infinity, especially in small and high Sérsic index profiles. To be safe, you can specify a large enough truncation radius for such compact high Sérsic index profiles.

If oversampling is used then the brightness is calculated using the over-sampled image, see Section 8.1.1.6 [Oversampling], page 200, which is much more accurate. The profile is first built in an array completely bounding it with a normalization constant of unity (see Section 8.1.1.4 [Galaxies], page 198). Taking B to be the desired brightness and S to be the sum of the pixels in the created profile, every pixel is then multiplied by B/S so the sum is exactly B.

If the `--individual` option is called, this same array is written to a FITS file. If not, only the overlapping pixels of this array and the output image are kept and added to the output array.

8.1.5 Invoking MakeProfiles

MakeProfiles will make any number of profiles specified in a catalog either individually or in one image. The executable name is `astmkprof` with the following general template

```
$ astmkprof [OPTION ...] [BackgroundImage] Catalog
```

One line examples:

```
## Make an image with profiles in catalog.txt (with default size):
$ astmkprof catalog.txt

## Make the profiles in catalog.txt over image.fits:
$ astmkprof --background=image.fits catalog.txt

## Make profiles in catalog, using RA and Dec in the given column:
$ astmkprof --racol=RA_CENTER --ycol=DEC_CENTER catalog.txt

## Make a 1500x1500 merged image (oversampled 500x500) image along
## with an individual image for all the profiles in catalog:
$ astmkprof --individual --oversample 3 -x500 -y500 catalog.txt
```

If mock images are to be made, a catalog (which stores the parameters for each mock profile) is mandatory. The catalog can be in the FITS ASCII, FITS binary format, or plain text formats (see Section 4.5 [Tables], page 66). The columns related to each paramter can be determined both by number, or by match/search criteria using the column names, units, or comments. with the options ending in `col`, see below. Without any file given to the `--background` option, Makeprofiles will make the fully zero-valued image and build the profiles on that (its size can be set with the `--naxis1` and `--naxis2` options, and its main WCS paramters can also be defined). Besides the main image containing all the profiles it

is also possible to build on individual images (only enclosing one full profile to its trucation radius) with the `--individual`.

If a data image file (see Section 4.1.1.1 [Arguments], page 49) is given, the pixels of that image are used as the background value for every pixel. The flux value of each profile pixel will be added to the pixel in that background value. In this case the values to all options relating to the output size and WCS will be ignored if specified (for example `--naxis1`, `--naxis2` and `--prepforconv`) on the command-line or in the configuration files. Note that `--oversample` will remain active even if a background image is specified.

Please see Section 2.2 [Sufi simulates a detection], page 17, for a very complete tutorial explaining how one could use MakeProfiles in conjunction with other Gnuastro's programs to make a complete simulated image of a mock galaxy.

8.1.5.1 MakeProfiles catalog

The catalog can be in the FITS ASCII, FITS binary format, or plain text formats (see Section 4.5 [Tables], page 66). Its columns can be ordered in any desired manner, you can specify which columns belong to which parameters using the set of options ending with `col` in Section 8.1.5.2 [MakeProfiles options], page 203. For example through the `--xcol` and `--rcol` options, you can specify the column that contains the X axis position of the profile center and the radial parameter for the profile. See Section 4.5.3 [Selecting table columns], page 71, for a thorough discussion on how to identify a column in the input catalog.

The value for the profile center in the catalog (in the `--xcol` and `--ycol` columns) can be a floating point number so the profile center can be on any sub-pixel position. Note that pixel positions in the FITS standard start from 1 and an integer is the pixel center. So a 2D image actually starts from the position (0.5, 0.5). When a `--background` image with WCS information is provided, you may also use RA and Dec to identify the center of each profile.

In MakeProfiles, profile centers do not have to be in (overlap with) the final image. Even if only one pixel of the profile within the truncation radius overlaps with the final image size, the profile is built and included in the final image image. Profiles that are completely out of the image will not be created. You can use the output log file to see which profiles were within the image.

If PSF profiles (Moffat or Gaussian, see Section 8.1.1.2 [Point Spread Function], page 196) are in the catalog and the profiles are to be built in one image (when `--individual` is not used), it is assumed they are the PSF(s) you want to convolve your created image with. So by default, they will not be built in the output image but as separate files. The sum of pixels of these separate files will also be set to unity (1) so you are ready to convolve, see Section 6.3.1.1 [Convolution process], page 118. As a summary, the position and magnitude of PSF profile will be ignored. This behaviour can be disabled with the `--psfinimg` option. If you want to create all the profiles separately (with `--individual`) and you want the sum of the PSF profile pixels to be unity, you have to set their magnitudes in the catalog to the zero-point magnitude and be sure that the central positions of the profiles don't have any fractional part (the PSF center has to be in the center of the pixel).

8.1.5.2 MakeProfiles options

The common options that are shared by Gnuastro programs, are fully explained in Section 4.1.2 [Common options], page 52, and are not repeated here. Since there are no

image inputs, the--hdu option is ignored. The options can be classified into the following categories: Output, Profiles, Catalog and WCS. Below each one is reviewed.

Output:

`-k STR`
`--background=STR`

A background image FITS file to build the profiles on. The extension that contains the image should be specified with the `--backhdu` option, see below. When a background image is specified, it will be used to derive all the information about the output image. Hence, the following options will be ignored: `--naxis1`, `--naxis2`, `--crpix1`, `--crpix2`, `--crval1`, `--crval2`, `--resolution`, `--oversample`, and data type (see `--type` in Section 4.1.2.1 [Input/Output options], page 52).

The image will act like a canvas to build the profiles on: profile pixel values will be summed with the background image pixel values. With the `--replace` option you can disable this behavior and replace the profile pixels with the background pixels. If you want to use all the image information above, except for the pixel values (you want to have a blank canvas to build the profiles on, based on an input image), you can call `--clearcanvas`, to set all the input image's pixels to zero before starting to build the profiles over it (this is done in memory after reading the input, so nothing will happen to your input file).

`-B STR/INT`
`--backhdu=STR/INT`

The header data unit (HDU) of the file given to

`-x INT`
`--naxis1=INT`

The number of pixels in the output image along the first FITS axis (horizontal when viewed in SAO ds9). This is before over-sampling. For example if you call MakeProfiles with `--naxis1=100 --oversample=5` (assuming no shift due for later convolution), then the final image size along the first axis will be 500. If a background image is specified, any possible value to this option is ignored.

`-y INT`
`--naxis2=INT`

The number of pixels in the output image along the second FITS axis (vertical when viewed in SAO ds9), see the explanation for `--naxis1`.

`-s INT`
`--oversample=INT`

The scale to over-sample the profiles and final image. If not an odd number, will be added by one, see Section 8.1.1.6 [Oversampling], page 200. Note that this `--oversample` will remain active even if an input image is specified. If your input catalog is based on the background image, be sure to set `--oversample=1`.

`--psfinimg`

Build the possibly existing PSF profiles (Moffat or Gaussian) in the catalog into the final image. By default they are built separately so you can convolve your images with them, thus their magnitude and positions are ignored. With

this option, they will be built in the final image like every other galaxy profile. To have a final PSF in your image, make a point profile where you want the PSF and after convolution it will be the PSF.

`-i`
`--individual`

If this option is called, each profile is created in a separate FITS file within the same directory as the output and the row number of the profile (starting from zero) in the name. The file for each row's profile will be in the same directory as the final combined image of all the profiles and will have the final image's name as a suffix. So for example if the final combined image is named `./out/fromcatalog.fits`, then the first profile that will be created with this option will be named `./out/0_fromcatalog.fits`.

Since each image only has one full profile out to the truncation radius the profile is centered and so, only the sub-pixel position of the profile center is important for the outputs of this option. The output will have an odd number of pixels. If there is no oversampling, the central pixel will contain the profile center. If the value to `--oversample` is larger than unity, then the profile center is on any of the central `--oversample`'d pixels depending on the fractional value of the profile center.

If the fractional value is larger than half, it is on the bottom half of the central region. This is due to the FITS definition of a real number position: The center of a pixel has fractional value 0.00 so each pixel contains these fractions: .5 – .75 – .00 (pixel center) – .25 – .5.

`-m`
`--nomerged`

Don't make a merged image. By default after making the profiles, they are added to a final image with sides of `--naxis1` and `--naxis2` if they overlap with it.

Profiles:

`-r`
`--numrandom`

The number of random points used in the central regions of the profile, see Section 8.1.1.5 [Sampling from a function], page 199.

`-e`
`--envseed`

Use the value to the `GSL_RNG_SEED` environment variable to generate the random Monte Carlo sampling distribution, see Section 8.1.1.5 [Sampling from a function], page 199, and Section 8.2.1.4 [Generating random numbers], page 212.

`-t FLT`
`--tolerance=FLT`

The tolerance to switch from Monte Carlo integration to the central pixel value, see Section 8.1.1.5 [Sampling from a function], page 199.

`-p`
`--tunitinp`
The truncation column of the catalog is in units of pixels. By default, the truncation column is considered to be in units of the radial parameters of the profile (`--rcol`). Read it as 't-unit-in-p' for 'truncation unit in pixels'.

`-X INT`
`--xshift=INT`
Shift all the profiles and enlarge the image along the first FITS axis, see n in Section 8.1.2 [If convolving afterwards], page 200. This is useful when you want to convolve the image afterwards. If you are using an external PSF, be sure to oversample it to the same scale used for creating the mock images. If a background image is specified, any possible value to this option is ignored.

`-Y INT`
`--yshift=INT`
Similar to `--xshift` for the second FITS axis.

`-c`
`--prepforconv`
Shift all the profiles and enlarge the image based on half the width of the first Moffat or Gaussian profile in the catalog, considering any possible oversampling see Section 8.1.2 [If convolving afterwards], page 200. `--prepforconv` is only checked and possibly activated if `--xshift` and `--yshift` are both zero (after reading the command-line and configuration files). If a background image is specified, any possible value to this option is ignored.

`--magatpeak`
The magnitude column in the catalog (see Section 8.1.5.1 [MakeProfiles catalog], page 203) will be used to find the brightness only for the peak profile pixel, not the full profile. Note that this is the flux of the profile's peak pixel in the final output of MakeProfiles. So beware of the oversampling, see Section 8.1.1.6 [Oversampling], page 200.

This option can be useful if you want to check a mock profile's total magnitude at various truncation radii. Without this option, no matter what the truncation radius is, the total magnitude will be the same as that given in the catalog. But with this option, the total magnitude will become brighter as you increase the truncation radius.

In sharper profiles, sometimes the accuracy of measuring the peak profile flux is more than the overall object brightness. In such cases, with this option, the final profile will be built such that its peak has the given magnitude, not the total profile.

> **CAUTION:** If you want to use this option for comparing with observations, please note that MakeProfiles does not do convolution. Unless you have deconvolved your data, your images are convolved with the instrument and atmospheric PSF, see Section 8.1.1.2 [Point Spread Function], page 196. Particularly in sharper profiles, the flux in the peak pixel is strongly decreased after convolution. Also note that in such cases, besides de-convolution, you will have to set `--oversample=1` otherwise after resampling your profile with Warp (see Section 6.4 [Warp], page 138), the peak flux will be different.

`-R`
`--replace`

Do not add the pixels of each profile over the background (possibly crowded by other profiles), replace them. By default, when two profiles overlap, the final pixel value is the sum of all the profiles that overlap on that pixel. When this option is given, the pixels are not added but replaced by the newer profile's pixel and any value under it is lost.

When order matters, make sure to use this function with '`--numthreads=1`'. When multiple threads are used, the separate profiles are built asynchronously and not in order. Since order does not matter in an addition, this causes no problems by default but has to be considered when this option is given. Using multiple threads is no problem if the profiles are to be used as a mask with a blank or fixed value (see '`--mforflatpix`') since all their pixel values are the same.

Note that only non-zero pixels are replaced. With radial profiles (for example Sérsic or Moffat) only values above zero will be part of the profile. However, when using flat profiles with the '`--mforflatpix`' option, you should be careful not to give a `0.0` value as the flat profile's pixel value.

`-C`

`--clearcanvas`

When an input image is specified (with the `--background` option, set all its pixels to 0.0 immediately after reading it into memory. Effectively, this will allow you to use all its properties (described under the `--background` option), without having to worry about the pixel values.

`--clearcanvas` can come in handy in many situations, for example if you want to create a labeled image (segmentation map) for creating a catalog (see Section 7.3 [MakeCatalog], page 175). In other cases, you might have modeled the objects in an image and want to create them on the same frame, but without the original pixel values.

`-w FLT`
`--circumwidth=FLT`

The width of the circumference if the profile is to be an elliptical circumference or annulus. See the explanations for this type of profile in `--fcol`.

`-z FLT`
`--zeropoint=FLT`

The zero-point magnitude of the image.

Catalog: The value to all of these options is considered to be a column number, where counting starts from zero.

`--fcol=INT/STR`

The functional form of the profile with one of the values below depending on the desired profile. The column can contain either the numeric codes (for example '`1`') or string characters (for example '`sersic`'). The numeric codes are easier to use in scripts which generate catalogs with hundreds or thousands of profiles.

The string format can be easier when the catalog is to be written/checked by hand/eye before running MakeProfiles. It is much more readable and provides a level of documentation. All Gnuastro's recognized table formats (see Section 4.5.1 [Recognized table formats], page 67) accept string type columns. To have string columns in a plain text table/catalog, see Section 4.5.2 [Gnuastro text table format], page 69.

- Sérsic profile with '`sersic`' or '`1`'.
- Moffat profile with '`moffat`' or '`2`'.
- Gaussian profile with '`gaussian`' or '`3`'.
- Point source with '`point`' or '`4`'.
- Flat profile with '`flat`' or '`5`'.
- Circumference profile with '`circum`' or '`6`'. A fixed value will be used for all pixels between the truncation radius (r_t) and $r_t - w$ (w is the value to the `--circumwidth`).

`--xcol=STR/INT`

The center of the profiles along the first FITS axis (horizontal when viewed in SAO ds9). See the explanations for `--racol` for precedence when both image and WCS coordinate columns are given.

`--ycol=STR/INT`

The center of the profiles along the second FITS axis (vertical when viewed in SAO ds9). Similar to `--xcol`.

`--racol=STR/INT`

The profile center's right ascension. Along with `--deccol`, these WCS coordinate columns are not mandatory. If they are not given, the `--xcol` and `--ycol` options will be used to specify the profile's central position and vice-versa. However, if image coordinate columns (`--xcol` and `--ycol`) and WCS coordinate columns (`--racol` and `--deccol`) are given, the WCS coordinate columns take precedence and image coordinate columns will be ignored.

`--deccol=STR/INT`

The profile center's declination. Similar to `--racol`.

`--rcol=STR/INT`

The radius parameter of the profiles. Effective radius (r_e) if Sérsic, FWHM if Moffat or Gaussian.

`--ncol=STR/INT`

The Sérsic index (n) or Moffat β.

`--pcol=STR/INT`
: The position angle (in degrees) of the profiles relative to the first FITS axis (horizontal when viewed in SAO ds9).

`--qcol=STR/INT`
: The axis ratio of the profiles (minor axis divided by the major axis).

`--mcol=STR/INT`
: The total pixelated magnitude of the profile within the truncation radius, see Section 8.1.4 [Profile magnitude], page 201.

`--tcol=STR/INT`
: The truncation radius of this profile. By default it is in units of the radial parameter of the profile (the value in the `--rcol` of the catalog). If `--tunitinp` is given, this value is interpreted in units of pixels (prior to oversampling) irrespective of the profile.

`-f`
`--mforflatpix`
: When making fixed value profiles (flat and circumference, see '`--fcol`'), don't use the value in the column specified by '`--mcol`' as the magnitude. Instead use it as the exact value that all the pixels of these profiles should have. This option is irrelevant for other types of profiles. This option is very useful for creating masks, or labeled regions in an image. Any integer, or floating point value can used in this column with this option, including `NaN` (or '`nan`', or '`NAN`', case is irrelevant), and infinities (`inf`, `-inf`, or `+inf`).

For example, with this option if you set the value in the magnitude column (`--mcol`) to `NaN`, you can create an elliptical or circular mask over an image (which can be given as the argument), see Section 6.1.3 [Blank pixels], page 100. Another useful application of this option is to create labeled elliptical or circular apertures in an image. To do this, set the value in the magnitude column to the label you want for this profile. This labeled image can then be used in combination with NoiseChisel's output (see Section 7.2.2.4 [NoiseChisel output], page 175) to do aperture photometry with MakeCatalog (see Section 7.3 [MakeCatalog], page 175).

Alternatively, if you want to mark regions of the image (for example with an elliptical circumference) and you don't want to use NaN values (as explained above) for some technical reason, you can get the minimum or maximum value in the image[9] using Arithmetic (see Section 6.2 [Arithmetic], page 108), then use that value in the magnitude column along with this option for all the profiles.

Please note that when using MakeProfiles on an already existing image, you have to set '`--oversample=1`'. Otherwise all the profiles will be scaled up based on the oversampling scale in your configuration files (see Section 4.2 [Configuration files], page 59) unless you have accounted for oversampling in your catalog.

[9] The minimum will give a better result, because the maximum can be too high compared to most pixels in the image, making it harder to display.

WCS:

`--crpix1=FLT`
: The pixel coordinates of the WCS reference point on the first (horizontal) FITS axis (counting from 1).

`--crpix2=FLT`
: The pixel coordinates of the WCS reference point on the second (vertical) FITS axis (counting from 1).

`--crval1=FLT`
: The Right Ascension (RA) of the reference point.

`--crval2=FLT`
: The Declination of the reference point.

`--resolution=FLT`
: The resolution of the non-oversampled image in units of arcseconds/pixel.

8.1.5.3 MakeProfiles output

Besides the final merged image of all the profiles or individual profiles that can be built based on the input options, MakeProfiles will also create a log file in the current directory (where you run MockProfiles). The values for each column are explained in the first few commented (starting with # character). The log file includes the following information:

- The total magnitude of the profile in the image. This will be different from your input magnitude if the profile was not completely in the image.
- The number of pixels (in the oversampled image) which used Monte Carlo integration and not the central pixel value.
- The fraction of flux in the Monte Carlo integrated pixels.
- If an individual image was created or not.

8.2 MakeNoise

Real data are always buried in noise, therefore to finalize a simulation of real data (for example to test our observational algorithms) it is essential to add noise to the mock profiles created with MakeProfiles, see Section 8.1 [MakeProfiles], page 195. Below, the general principles and concepts to help understand how noise is quantified is discussed. MakeNoise options and argument are then discussed in Section 8.2.2 [Invoking MakeNoise], page 214.

8.2.1 Noise basics

Deep astronomical images, like those used in extragalactic studies seriously suffer from noise in the data. Generally speaking, the sources of noise in an astronomical image are photon counting noise and Instrumental noise which are discussed in detail below. We finish with a short introduction on how random numbers are generated and how you can determine the random number generator and seed value.

8.2.1.1 Photon counting noise

Thanks to the very accurate electronics used in today's detectors, this type of noise is the main cause of concern for extra galactic studies. It can generally be associate with the

counting error that is known to have a Poisson distribution. The Poisson distribution is about counting. But counting is a discrete operation with only positive values, for example we can't count 3.2 or -2 of anything. We only count 0, 1, 2, 3 and so on. Therefore the Poisson distribution is also a discrete distribution, only applying to whole positive integers.

Let's assume the mean value of counting something is known. In this case, the number of electrons that are produced by photons in the CCD. Let's call this mean λ. Let's take k to represent the result of counting in one particular time we attempt to count. The probability density function of k can be written as:

$$f(k) = \frac{\lambda^k}{k!}e^{-\lambda}, \quad k \in \{0, 1, 2, 3, \ldots\}$$

Because the Poisson distribution is only applicable to positive values, it is by nature very skewed when λ is near zero. One qualitative way to imagine it is that there simply aren't enough integers smaller than λ, than there are larger integers. Therefore to accommodate all possibilities, it has to be skewed when λ is small.

But as λ becomes larger and larger, the distribution becomes more and more symmetric. One very useful property of the Poisson distribution is that the mean value is also its variance. When λ is very large, say $\lambda > 1000$, then the normal (Gaussian) distribution, see Section 8.1.1.2 [Point Spread Function], page 196, is an excellent approximation of the Poisson distribution with mean $\mu = \lambda$ and standard deviation $\sigma = \sqrt{\lambda}$.

We see that the variance or dispersion of the distribution depends on the mean value, and when it is large it can be approximated with a Gaussian that only has one free parameter ($\mu = \lambda$ and $\sigma = \sqrt{\lambda}$) instead of two that it originally has.

The astronomical objects after convolution with the PSF of the instrument, lie above a certain background flux. This background flux is defined to be the average flux of a region in the image that has absolutely no objects. The physical origin of this background value is the brightness of the atmosphere or possible stray light within the imaging instrument. It is thus an ideal definition, because in practice, what lies deep in the noise far lower than the detection limit is never known[10]. However, in a real image, a relatively large number of very faint objects can been fully buried in the noise. These undetected objects will bias the background measurement to slightly larger values. The sky value is therefore defined to be the average of the undetected regions in the image, so in an ideal case where all the objects have been detected, the sky value and background value are the same.

As longer wavelengths are used, the background value becomes more significant and also varies over a wide image field. Such variations are not currently implemented in MakeProfiles, but will be in the future. In a mock image, we have the luxury of setting the background value.

In each pixel of the canvas of pixels, the flux is the sum of contributions from various sources after convolution. Let's name this flux of the convolved sum of possibly overlapping objects, I_{nn}. nn representing 'no noise'. For now, let's assume the background is constant and represented by B. In practice the background values are larger than $\sim 1,000$ counts. Then the flux after adding noise is a random value taken from a Gaussian distribution with the following mean (μ) and standard deviation (σ):

[10] See the section on sky in Akhlaghi M., Ichikawa. T. 2015. Astrophysical Journal Supplement Series.

$$\mu = B + I_{nn}, \quad \sigma = \sqrt{B + I_{nn}}$$

Since this type of noise is inherent in the objects we study, it is usually measured on the same scale as the astronomical objects, namely the magnitude system, see Section 8.1.3 [Flux Brightness and magnitude], page 201. It is then internally converted to the flux scale for further processing.

8.2.1.2 Instrumental noise

While taking images with a camera, a dark current is fed to the pixels, the variation of the value of this dark current over the pixels, also adds to the final image noise. Another source of noise is the readout noise that is produced by the electronics in the CCD that attempt to digitize the voltage produced by the photo-electrons in the analog to digital converter. In deep extra-galactic studies these sources of noise are not as significant as the noise of the background sky. Let C represent the combined standard deviation of all these sources of noise. If only this source of noise is present, the noised pixel value would be a random value chosen from a Gaussian distribution with

$$\mu = I_{nn}, \quad \sigma = \sqrt{C^2 + I_{nn}}$$

This type of noise is completely independent of the type of objects being studied, it is completely determined by the instrument. So the flux scale (and not magnitude scale) is most commonly used for this type of noise. In practice, this value is usually reported in ADUs not flux or electron counts. The gain value of the device can be used to convert between these two, see Section 8.1.3 [Flux Brightness and magnitude], page 201.

8.2.1.3 Final noised pixel value

Depending on the values you specify for B and C from the above, the final noised value for each pixel is a random value chosen from a Gaussian distribution with

$$\mu = B + I_{nn}, \quad \sigma = \sqrt{C^2 + B + I_{nn}}$$

8.2.1.4 Generating random numbers

As discussed above, to generate noise we need to make random samples of a particular distribution. So it is important to understand some general concepts regarding the generation of random numbers. For a very complete and nice introduction we strongly advise reading Donald Knuth's "The art of computer programming", volume 2, chapter 3[11]. Quoting from the GNU Scientific Library manual, "If you don't own it, you should stop reading right now, run to the nearest bookstore, and buy it"[12]!

[11] Knuth, Donald. 1998. The art of computer programming. Addison–Wesley. ISBN 0-201-89684-2

[12] For students, running to the library might be more affordable!

Using only software, we can only produce what is called a psudo-random sequence of numbers. A true random number generator is a hardware (let's assume we have made sure it has no systematic biases), for example throwing dice or flipping coins (which have remained from the ancient times). More modern hardware methods use atmospheric noise, thermal noise or other types of external electromagnetic or quantum phenomena. All psudo-random number generators (software) require a seed to be the basis of the generation. The advantage of having a seed is that if you specify the same seed for multiple runs, you will get an identical sequence of random numbers which allows you to reproduce the same final noised image.

The programs in GNU Astronomy Utilities (for example MakeNoise or MakeProfiles) use the GNU Scientific Library (GSL) to generate random numbers. GSL allows the user to set the random number generator through environment variables, see Section 3.3.1.2 [Installation directory], page 39, for an introduction to environment variables. In the chapter titled "Random Number Generation" they have fully explained the various random number generators that are available (there are a lot of them!). Through the two environment variables `GSL_RNG_TYPE` and `GSL_RNG_SEED` you can specify the generator and its seed respectively.

If you don't specify a value for `GSL_RNG_TYPE`, GSL will use its default random number generator type. The default type is sufficient for most general applications. If no value is given for the `GSL_RNG_SEED` environment variable and you have asked Gnuastro to read the seed from the environment (through the `--envseed` option), then GSL will use the default value of each generator to give identical outputs. If you don't explicitly tell Gnuastro programs to read the seed value from the environment variable, then they will use the system time (accurate to within a microsecond) to generate (apparently random) seeds. In this manner, every time you run the program, you will get a different random number distribution.

There are two ways you can specify values for these environment variables. You can call them on the same command-line for example:

```
$ GSL_RNG_TYPE="taus" GSL_RNG_SEED=345 astmknoise input.fits
```

In this manner the values will only be used for this particular execution of MakeNoise. Alternatively, you can define them for the full period of your terminal session or script length, using the shell's `export` command with the two separate commands below (for a script remove the `$` signs):

```
$ export GSL_RNG_TYPE="taus"
$ export GSL_RNG_SEED=345
```

The subsequent programs which use GSL's random number generators will hence forth use these values in this session of the terminal you are running or while executing this script. In case you want to set fixed values for these parameters every time you use the GSL random number generator, you can add these two lines to your `.bashrc` startup script[13], see Section 3.3.1.2 [Installation directory], page 39.

[13] Don't forget that if you are going to give your scripts (that use the GSL random number generator) to others you have to make sure you also tell them to set these environment variable separately. So for scripts, it is best to keep all such variable definitions within the script, even if they are within your `.bashrc`.

NOTE: If the two environment variables `GSL_RNG_TYPE` and `GSL_RNG_SEED` are defined, GSL will report them by default, even if you don't use the `--envseed` option. For example you can see the top few lines of the output of MakeProfiles:

```
$ export GSL_RNG_TYPE="taus"
$ export GSL_RNG_SEED=345
$ astmkprof catalog.txt --envseed
GSL_RNG_TYPE=taus
GSL_RNG_SEED=345
MakeProfiles started on AAA BBB DD EE:FF:GG HHH
  - 6 profiles read from catalog.txt 0.000236 seconds
  - Random number generator (RNG) type: taus
  - RNG seed for all profiles: 345
```

The first two output lines (showing the names of the environment variables) are printed by GSL before MakeProfiles actually starts generating random numbers. The Gnuastro programs will report the values they use independently, you should check them for the final values used. For example if `--envseed` is not given, `GSL_RNG_SEED` will not be used and the last line shown above will not be printed. In the case of MakeProfiles, each profile will get its own seed value.

8.2.2 Invoking MakeNoise

MakeNoise will add noise to an existing image. The executable name is `astmknoise` with the following general template

```
$ astmknoise [OPTION ...] InputImage.fits
```

One line examples:

```
## Add noise to input image assuming background and instrumental noise:
$ astmknoise --background=1000 --stdadd=20 mockimage.fits
```

If actual processing is to be done, the input image is a mandatory argument. The full list of options common to all the programs in Gnuastro can be seen in Section 4.1.2 [Common options], page 52. The output will have the same type as the input image, however the internal processing is done on a double precision floating point format. If the input values were integer types, then each floating point number will be rounded to the nearest integer away from zero. This might cause integer overflow if types with small ranges are used (for example images with a `BITPIX` of `8` which can only keep 256 values). This can be disabled with the `doubletype` option. The header of the output FITS file keeps all the parameters that were influential in making it. This is done for future reproducibility.

`-b FLT`
`--background=FLT`

The background pixel value for the image in units of magnitudes, see Section 8.2.1.1 [Photon counting noise], page 210, and Section 8.1.3 [Flux Brightness and magnitude], page 201.

`-z FLT`
`--zeropoint=FLT`

The zeropoint magnitude used to convert the value of `--background` (in units of magnitude) to flux, see Section 8.1.3 [Flux Brightness and magnitude], page 201.

`-s FLT`
`--stdadd=FLT`
: The instrumental noise which is in units of flux, see Section 8.2.1.2 [Instrumental noise], page 212.

`-e`
`--envseed`
: Use the `GSL_RNG_SEED` environment variable for the seed used in the random number generator, see Section 8.2.1.4 [Generating random numbers], page 212. With this option, the output image noise is always going to be identical (or reproducible).

`-d`
`--doubletype`
: Save the output in the double precision floating point format that was used internally. This option will be most useful if the input images were of integer types.

9 High-level calculations

After the reduction of raw data (for example with the programs in Chapter 6 [Data manipulation], page 97) you will have reduced images/data ready for processing/analyzing (for example with the programs in Chapter 7 [Data analysis], page 148). But the processed/analyzed data (or catalogs) are still not enough to derive any scientific result. Even higher-level analysis is still needed to convert the observed magnitudes, sizes or volumes into physical quantities that we associate with each catalog entry or detected object which is the purpose of the tools in this section.

9.1 CosmicCalculator

To derive higher-level information regarding our sources in extra-galactic astronomy, cosmological calculations are necessary. In Gnuastro, CosmicCalculator is in charge of such calculations. Before discussing how CosmicCalculator is called and operates (in Section 9.1.3 [Invoking CosmicCalculator], page 221), it is important to provide a rough but mostly self sufficient review of the basics and the equations used in the analysis. In Section 9.1.1 [Distance on a 2D curved space], page 216, the basic idea of understanding distances in a curved and expanding 2D universe (which we can visualize) are reviewed. Having solidified the concepts there, in Section 9.1.2 [Extending distance concepts to 3D], page 220, the formalism is extended to the 3D universe we are trying to study in our research.

The focus here is obtaining a physical insight into these equations (mainly for the use in real observational studies). There are many books thoroughly deriving and proving all the equations with all possible initial conditions and assumptions for any abstract universe, interested readers can study those books.

9.1.1 Distance on a 2D curved space

The observations to date (for example the Plank 2013 results), have not measured the presence of a significant curvature in the universe. However to be generic (and allow its measurement if it does in fact exist), it is very important to create a framework that allows curvature. As 3D beings, it is impossible for us to mentally create (visualize) a picture of the curvature of a 3D volume in a 4D space. Hence, here we will assume a 2D surface and discuss distances on that 2D surface when it is flat, or when the 2D surface is curved (in a 3D space). Once the concepts have been created/visualized here, in Section 9.1.2 [Extending distance concepts to 3D], page 220, we will extend them to the real 3D universe we live in and hope to study.

To be more understandable (actively discuss from an observer's point of view) let's assume we have an imaginary 2D friend living on the 2D space (which *might* be curved in 3D). So here we will be working with it in its efforts to analyze distances on its 2D universe. The start of the analysis might seem too mundane, but since it is impossible to imagine a 3D curved space, it is important to review all the very basic concepts thoroughly for an easy transition to a universe we cannot visualize any more (a curved 3D space in 4D).

To start, let's assume a static (not expanding or shrinking), flat 2D surface similar to Figure 9.1 and that our 2D friend is observing its universe from point A. One of the most basic ways to parametrize this space is through the Cartesian coordinates (x, y). In Figure 9.1, the basic axes of these two coordinates are plotted. An infinitesimal change in

the direction of each axis is written as dx and dy. For each point, the infinitesimal changes are parallel with the respective axes and are not shown for clarity. Another very useful way of parametrizing this space is through polar coordinates. For each point, we define a radius (r) and angle (ϕ) from a fixed (but arbitrary) reference axis. In Figure 9.1 the infinitesimal changes for each polar coordinate are plotted for a random point and a dashed circle is shown for all points with the same radius.

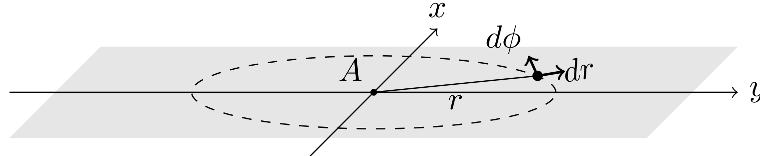

Figure 9.1: Two dimensional Cartesian and polar coordinates on a flat plane.

Assuming a certain position, which can be parametrized as (x, y), or (r, ϕ), a general infinitesimal change change in its position will place it in the coordinates $(x+dx, y+dy)$ and $(r + dr, \phi + d\phi)$. The distance (on the flat 2D surface) that is covered by this infinitesimal change in the static universe (ds_s, the subscript signifies the static nature of this universe) can be written as

$$ds_s = dx^2 + dy^2 = dr^2 + r^2 d\phi^2$$

The main question is this: how can our 2D friend incorporate the (possible) curvature in its universe when it is calculating distances? The universe it lives in might equally be a locally flat but globally curved surface like Figure 9.2. The answer to this question but for a 3D being (us) is the whole purpose to this discussion. So here we want to give our 2D friend (and later, ourselves) the tools to measure distances if the space (that hosts the objects) is curved.

Figure 9.2 assumes a spherical shell with radius R as the curved 2D plane for simplicity. The spherical shell is tangent to the 2D plane and only touches it at A. The result will be generalized afterwards. The first step in measuring the distance in a curved space is to imagine a third dimension along the z axis as shown in Figure 9.2. For simplicity, the z axis is assumed to pass through the center of the spherical shell. Our imaginary 2D friend cannot visualize the third dimension or a curved 2D surface within it, so the remainder of this discussion is purely abstract for it (similar to us being unable to visualize a 3D curved space in 4D). But since we are 3D creatures, we have the advantage of visualizing the following steps. Fortunately our 2D friend knows our mathematics, so it can follow along with us.

With the third axis added, a generic infinitesimal change over *the full* 3D space corresponds to the distance:

$$ds_s^2 = dx^2 + dy^2 + dz^2 = dr^2 + r^2 d\phi^2 + dz^2.$$

It is very important to recognize that this change of distance is for *any* point in the 3D space, not just those changes that occur on the 2D spherical shell of Figure 9.2. Recall that our 2D friend can only do measurements in the 2D spherical shell, not the full 3D space. So we have to constrain this general change to any change on the 2D spherical shell. To

do that, let's look at the arbitrary point P on the 2D spherical shell. Its image (P') on the flat plain is also displayed. From the dark triangle, we see that

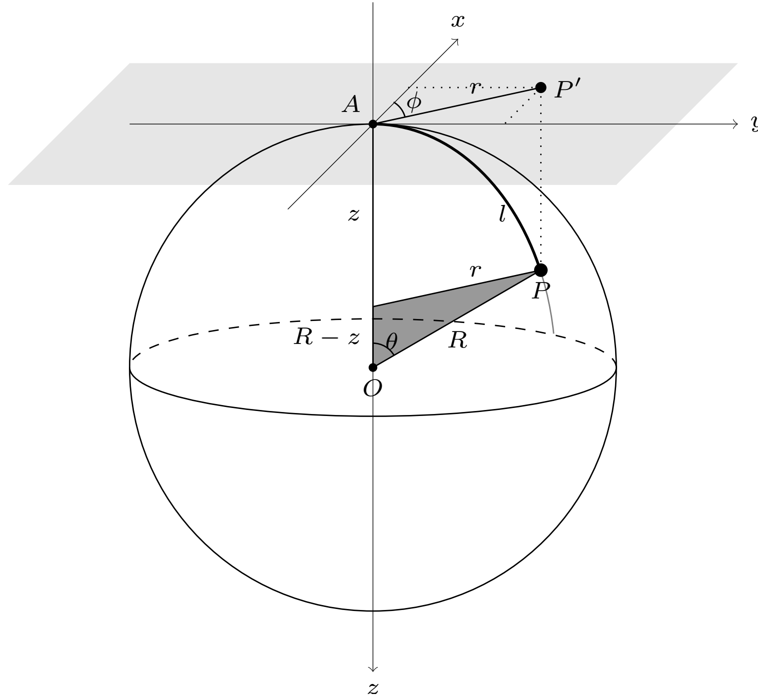

Figure 9.2: 2D spherical plane (centered on O) and flat plane (gray) tangent to it at point A.

$$\sin\theta = \frac{r}{R}, \quad \cos\theta = \frac{R-z}{R}.$$

These relations allow our 2D friend to find the value of z (an abstract dimension for it) as a function of r (distance on a flat 2D plane, which it can visualize) and thus eliminate z. From $\sin^2\theta + \cos^2\theta = 1$, we get $z^2 - 2Rz + r^2 = 0$ and solving for z, we find:

$$z = R\left(1 \pm \sqrt{1 - \frac{r^2}{R^2}}\right).$$

The $\pm$ can be understood from Figure 9.2: For each r, there are two points on the sphere, one in the upper hemisphere and one in the lower hemisphere. An infinitesimal change in r, will create the following infinitesimal change in z:

$$dz = \frac{\mp r}{R}\left(\frac{1}{\sqrt{1 - r^2/R^2}}\right) dr.$$

Using the positive signed equation instead of dz in the ds_s^2 equation above, we get:

$$ds_s^2 = \frac{dr^2}{1 - r^2/R^2} + r^2 d\phi^2.$$

The derivation above was done for a spherical shell of radius R as a curved 2D surface. To generalize it to any surface, we can define $K = 1/R^2$ as the curvature parameter. Then the general infinitesimal change in a static universe can be written as:

$$ds_s^2 = \frac{dr^2}{1 - Kr^2} + r^2 d\phi^2.$$

Therefore, we see that a positive K represents a real R which signifies a closed 2D spherical shell like Figure 9.2. When $K = 0$, we have a flat plane (Figure 9.1) and a negative K will correspond to an imaginary R. The latter two cases are open universes (where r can extend to infinity). However, when $K > 0$, we have a closed universe, where r cannot become larger than R as in Figure 9.2.

A very important issue that can be discussed now (while we are still in 2D and can actually visualize things) is that $\overrightarrow{r}$ is tangent to the curved space at the observer's position. In other words, it is on the gray flat surface of Figure 9.2, even when the universe if curved: $\overrightarrow{r} = P' - A$. Therefore for the point P on a curved space, the raw coordinate r is the distance to P', not P. The distance to the point P (at a specific coordinate r on the flat plane) on the curved surface (thick line in Figure 9.2) is called the *proper distance* and is displayed with l. For the specific example of Figure 9.2, the proper distance can be calculated with: $l = R\theta$ (θ is in radians). using the $\sin\theta$ relation found above, we can find l as a function of r:

$$\theta = \sin^{-1}\left(\frac{r}{R}\right) \quad \rightarrow \quad l(r) = R\sin^{-1}\left(\frac{r}{R}\right)$$

R is just an arbitrary constant and can be directly found from K, so for cleaner equations, it is common practice to set $R = 1$, which gives: $l(r) = \sin^{-1} r$. Also note that if $R = 1$, then $l = \theta$. Generally, depending on the the curvature, in a *static* universe the proper distance can be written as a function of the coordinate r as (from now on we are assuming $R = 1$):

$$l(r) = \sin^{-1}(r) \quad (K > 0), \qquad l(r) = r \quad (K = 0), \qquad l(r) = \sinh^{-1}(r) \quad (K < 0).$$

With l, the infinitesimal change of distance can be written in a more simpler and abstract form of

$$ds_s^2 = dl^2 + r^2 d\phi^2.$$

Until now, we had assumed a static universe (not changing with time). But our observations so far appear to indicate that the universe is expanding (isn't static). Since there is no reason to expect the observed expansion is unique to our particular position of the universe, we expect the universe to be expanding at all points with the same rate at the same time. Therefore, to add a time dependence to our distance measurements, we can simply add a multiplicative scaling factor, which is a function of time: $a(t)$. The functional form of $a(t)$ comes from the cosmology and the physics we assume for it: general relativity.

With this scaling factor, the proper distance will also depend on time. As the universe expands (moves), the distance will also move to larger values. We thus define a distance measure, or coordinate, that is independent of time and thus doesn't 'move' which we call the *comoving distance* and display with χ such that: $l(r,t)=\chi(r)a(t)$. We thus shift the r dependence of the proper distance we derived above for a static universe to the comoving distance:

$$\chi(r)=\sin^{-1}(r)\quad (K>0),\qquad \chi(r)=r\quad (K=0),\qquad \chi(r)=\sinh^{-1}(r)\quad (K<0).$$

Therefore $\chi(r)$ is the proper distance of an object at a specific reference time: $t=t_r$ (the r subscript signifies "reference") when $a(t_r)=1$. At any arbitrary moment ($t\neq t_r$) before or after t_r, the proper distance to the object can simply be scaled with $a(t)$. Measuring the change of distance in a time-dependent (expanding) universe will also involve the speed of the object changing positions. Hence, let's assume that we are only thinking about the change in distance caused by something (light) moving at the speed of light. This speed is postulated as the only constant and frame-of-reference-independent speed in the universe, making our calculations easier, light is also the major source of information we receive from the universe, so this is a reasonable assumption for most extra-galactic studies. We can thus parametrize the change in distance as

$$ds^2=c^2dt^2-a^2(t)ds_s^2=c^2dt^2-a^2(t)(d\chi^2+r^2d\phi^2).$$

9.1.2 Extending distance concepts to 3D

The concepts of Section 9.1.1 [Distance on a 2D curved space], page 216, are here extended to a 3D space that *might* be curved in a 4D space. We can start with the generic infinitesimal distance in a static 3D universe, but this time not in spherical coordinates instead of polar coordinates. θ is shown in Figure 9.2, but here we are 3D beings, positioned on O (the center of the sphere) and the point O is tangent to a 4D-sphere. In our 3D space, a generic infinitesimal displacement will have the distance:

$$ds_s^2=dx^2+dy^2+dz^2=dr^2+r^2(d\theta^2+\sin^2\theta d\phi^2).$$

Like our 2D friend before, we now have to assume an abstract dimension which we cannot visualize. Let's call the fourth dimension w, then the general change in coordinates in the *full* four dimensional space will be:

$$ds_s^2=dr^2+r^2(d\theta^2+\sin^2\theta d\phi^2)+dw^2.$$

But we can only work on a 3D curved space, so following exactly the same steps and conventions as our 2D friend, we arrive at:

$$ds_s^2=\frac{dr^2}{1-Kr^2}+r^2(d\theta^2+\sin^2\theta d\phi^2).$$

In a non-static universe (with a scale factor a(t), the distance can be written as:

$$ds^2=c^2dt^2-a^2(t)[d\chi^2+r^2(d\theta^2+\sin^2\theta d\phi^2)].$$

9.1.3 Invoking CosmicCalculator

CosmicCalculator will calculate cosmological variables based on the input parameters. The executable name is `astcosmiccal` with the following general template

```
$ astcosmiccal [OPTION...] ...
```

One line examples:

```
## Print basic cosmological properties at redshift 2.5:
$ astcosmiccal -z2.5

## Only print Comoving volume over 4pi stradian to z (Mpc^3):
$ astcosmiccal --onlyvolume --redshift=0.8

## Assume Lambda and matter density of 0.7 and 0.3 and print
## basic cosmological parameters for redshift 2.1:
$ astcosmiccal -l0.7 -m0.3 -z2.1
```

The input parameters can be given as command-line options or in the configuration files, see Section 4.2 [Configuration files], page 59. For a definition of the different parameters, please see the sections prior to this. By default, all the cosmological calculations will be printed in the standard output (the command-line mainly) along with a short description and units.

The options starting with `--only` will only do that single desired calculation and only print the final number (in the same units as reported by default). These options are very useful when you want to call CosmicCalculator from a script. The resulting number can simply be put into a shell variable (for example `vol`) with the following line, which will allow you to use the value for any other subsequent operation.

```
z=3.12
vol=$(astcosmiccal --redshift=$z --onlyvolume)
```

In a script, this operation might be necessary for a very large number of objects (thousands of galaxies in a catalog for example). So the fact that all the other default calculations are ignored will also help you get to your result faster. If you just want to inspect the value of a variable, the description (which comes with units) might be more useful. In that case, the following command might be better. The other parameters will also be calculated, but they are so fast that you will not notice on modern computers.

```
$ astcosmiccal --redshift=0.832 | grep volume
```

The full list of options is shown and described below:

`-z FLT`
`--redshift=FLT`
The redshift of interest.

`-H FLT`
`--H0=FLT` Current expansion rate (in km sec^{-1} Mpc^{-1}).

`-l FLT`
`--olambda=FLT`
Cosmological constant density divided by the critical density in the current Universe ($\Omega_{\Lambda,0}$).

`-m FLT`
`--omatter=FLT`
: Matter (including massive neutrinos) density divided by the critical density in the current Universe ($\Omega_{m,0}$).

`-r FLT`
`--oradiation=FLT`
: Radiation density divided by the critical density in the current Universe ($\Omega_{r,0}$).

`-v`
`--onlyvolume`
: Only print the comoving volume (in units of Mpc^3) until the desired redshift based on the input parameters. See explanations above for more on these types of options and how to effectively use them.

`-d`
`--onlyabsmagconv`
: Only print the conversion factor for apparent magnitude to absolute magnitude. Note that this is practically the distance modulus added with $-2.5\log{(1+z)}$ for the the desired redshift based on the input parameters. See explanations above for more on these types of options and how to effectively use them.

10 Library

Each program in Gnuastro that was discussed in the prior chapters (or any program in general) is a collection of functions that is compiled into one executable file which can communicate directly with the outside world. The outside world in this context is the operating system. By communication, we mean that control is directly passed to a program from the operating system with a (possible) set of inputs and after it is finished, the program will pass control back to the operating system. For programs written in C and C++, the unique `main` function is in charge of this communication.

Similar to a program, a library is also a collection of functions that is compiled into one executable file. However, unlike programs, libraries don't have a `main` function. Therefore they can't communicate directly with the outside world. This gives you the chance to write your own `main` function and call library functions from within it. After compiling your program into a binary executable, you just have to *link* it to the library and you are ready to run (execute) your program. In this way, you can use Gnuastro at a much lower-level, and in combination with other libraries on your system, you can significantly boost your creativity.

This chapter starts with a basic introduction to libraries and how you can use them in Section 10.1 [Review of library fundamentals], page 223. The separate functions in the Gnuastro library are then introduced (classified by context) in Section 10.3 [Gnuastro library], page 233. If you end up routinely using a fixed set of library functions, with a well-defined input and output, it will be much more beneficial if you define a program for the job. Therefore, in its Section 3.2.2 [Version controlled source], page 32, Gnuastro comes with the Section 11.4.2 [The TEMPLATE program], page 329, to easily define your own programs(s).

10.1 Review of library fundamentals

Gnuastro's libraries are written in the C programming language. In Section 11.1 [Why C programming language?], page 320, we have thoroughly discussed the reasons behind this choice. C was actually created to write Unix, thus understanding the way C works can greatly help in effectively using programs and libraries in all Unix-like operating systems. Therefore, in the following subsections some important aspects of C, as it relates to libraries (and thus programs that depend on them) on Unix are reviewed. First we will discuss header files in Section 10.1.1 [Headers], page 224, and then go onto Section 10.1.2 [Linking], page 227. This section finishes with Section 10.1.3 [Summary and example on libraries], page 229. If you are already familiar with these concepts, please skip this section and go directly to Section 10.3 [Gnuastro library], page 233.

In theory, a full operating system (or any software) can be written as one function. Such a software would not need any headers or linking (that are discussed in the subsections below). However, writing that single function and maintaining it (adding new features, fixing bugs, documentation and etc) would be a programmer or scientist's worst nightmare! Furthermore, all the hard work that went into creating it cannot be reused in other software: every other programmer or scientist would have to re-invent the wheel. The ultimate purpose behind libraries (which come with headers and have to be linked) is to address this problem and increase modularity: "the degree to which a system's components may

be separated and recombined" (from Wikipedia). The more modular the source code of a program or library, the easier maintaining it will be, and all the hard work that went into creating it can be reused for a wider range of problems.

10.1.1 Headers

C source code is read from top to bottom in the source file, therefore program components (for example variables, data structures and functions) should all be *defined* or *declared* closer to the top of the source file: before they are used. *Defining* something in C or C++ is jargon for providing its full details. *Declaring* it, on the other-hand, is jargon for only providing the minimum information needed for the compiler to pass it temporarily and fill in the detailed definition later.

For a function, the *declaration* only contains the inputs and their data-types along with the output's type[1]. The *definition* adds to the declaration by including the exact details of what operations are done to the inputs to generate the output. As an example, take this simple summation function:

```
double
sum(double a, double b)
{
  return a + b;
}
```

What you see above is the *definition* of this function: it shows you (and the compiler) exactly what it does to the two `double` type inputs and that the output also has a `double` type. Note that a function's internal operations are rarely so simple and short, it can be arbitrarily long and complicated. This unreasonably short and simple function was chosen here for ease of reading. The declaration for this function is:

```
double
sum(double a, double b);
```

You can think of a function's declaration as a building's address in the city, and the definition as the building's complete blueprints. When the compiler confronts a call to a function during its processing, it doesn't need to know anything about how the inputs are processed to generate the output. Just as the postman doesn't need to know the inner structure of a building when delivering the mail. The declaration (address) is enough. Therefore by *declaring* the functions once at the start of the source files, we don't have to worry about *defining* them after they are used.

Even for a simple real-world operation (not a simple summation like above!), you will soon need many functions (for example, some for reading/preparing the inputs, some for the processing, and some for preparing the output). Although it is technically possible, managing all the necessary functions in one file is not easy and is contrary to the modularity principle (see Section 10.1 [Review of library fundamentals], page 223), for example the functions for preparing the input can be usable in your other projects with a different processing. Therefore, as we will see later (in Section 10.1.2 [Linking], page 227), the functions don't necessarily need to be defined in the source file where they are used. As long as their definitions are ultimately linked to the final executable, everything will be fine. For now, it is just important to remember that the functions that are called within

[1] Recall that in C, functions only have one output.

one source file must be declared within the source file (declarations are mandatory), but not necessarily defined there.

In the spirit of modularity, it is common to define contextually similar functions in one source file. For example, in Gnuastro, functions that calculate the median, mean and other statistical functions are defined in `lib/statistics.c`, while functions that deal directly with FITS files are defined in `lib/fits.c`.

Keeping the definition of similar functions in a separate file greatly helps their management and modularity, but this fact alone doesn't make things much easier for the caller's source code: recall that while definitions are optional, declarations are mandatory. So if this was all, the caller would have to manually copy and paste (*include*) all the declarations from the various source files into the file they are working on now. To address this problem, programmers have adopted the header file convention: the header file of a source code contains all the declarations that a caller would need to be able to use any of its functions. For example, in Gnuastro, `lib/statistics.c` (file containing function definitions) comes with `lib/gnuastro/statistics.h` (only containing function declarations).

The discussion above was mainly focused on functions, however, there are many more programming constructs such as pre-processor macros and data structures. Like functions, they also need to be known to the compiler when it confronts a call to them. So the header file also contains their definitions or declarations when they are necessary for the functions.

Pre-processor macros (or macros for short) are replaced with their defined value by the pre-processor before compilation. Conventionally they are written only in capital letters to be easily recognized. It is just important to understand that the compiler doesn't see the macros, it sees their fixed values. So when a header specifies macros you can do your programming without worrying about the actual values. The standard C types (for example `int`, or `float`) are very low-level and basic. We can collect multiple C types into a *structure* for a higher-level way to keep and pass-along data. See Section 10.3.5.1 [Generic data container (`gal_data_t`)], page 245, for some examples of macros and data structures.

The contents in the header need to be *include*d into the caller's source code with a special pre-processor command: `#include <path/to/header.h>`. As the name suggests, the *pre-processor* goes through the source code prior to the processor (or compiler). One of its jobs is to include, or merge, the contents of files that are mentioned with this directive in the source code. Therefore the compiler sees a single entity containing the contents of the main file and all the included files. This allows you to include many (sometimes thousands of) declarations into your code with only one line. Since the headers are also installed with the library into your system, you don't even need to keep a copy of them for each separate program, making things even more convenient.

Try opening some of the `.c` files in Gnuastro's `lib/` directory with a text editor to check out the include directives at the start of the file (after the copyright notice). Let's take `lib/fits.c` as an example. You will notice that Gnuastro's header files (like `gnuastro/fits.h`) are indeed within this directory (the `fits.h` file is in the `gnuastro/` directory). You will notice that files like `stdio.h`, or `string.h` are not in this directory (or anywhere within Gnuastro).

On most systems the basic C header files (like `stdio.h` and `string.h` mentioned above) are located in `/usr/include/`[2]. Your compiler is configured to automatically search that directory (and possibly others), so you don't have to explicitly mention these directories. Go ahead, look into the `/usr/include` directory and find `stdio.h` for example. When the necessary header files are not in those specific libraries, the pre-processor can also search in places other than the current directory. You can specify those directories with this pre-processor option[3]:

`-I DIR` "Add the directory `DIR` to the list of directories to be searched for header files. Directories named by '-I' are searched before the standard system include directories. If the directory `DIR` is a standard system include directory, the option is ignored to ensure that the default search order for system directories and the special treatment of system headers are not defeated..." (quoted from the GNU Compiler Collection manual). Note that the space between `I` and the directory is optional and commonly not used.

If the pre-processor can't find the included files, it will abort with an error. In fact a common error when building programs that depend on a library is that the compiler doesn't not know where a library's header is (see Section 3.3.4 [Known issues], page 46). So you have to manually tell the compiler where to look for the library's headers with the `-I` option. For a small software with one or two source files, this can be done manually (see Section 10.1.3 [Summary and example on libraries], page 229). However, to enhance modularity, Gnuastro (and most other bin/libraries) contain many source files, so the compiler is invoked many times[4]. This makes manual addition or modification of this option practically impossible.

To solve this problem, in the GNU build system, there are conventional environment variables for the various kinds of compiler options (or flags). These environment variables are used in every call to the compiler (they can be empty). The environment variable used for the C Pre-Processor (or CPP) is `CPPFLAGS`. By giving `CPPFLAGS` a value once, you can be sure that each call to the compiler will be affected. See Section 3.3.4 [Known issues], page 46, for an example of how to set this variable at configure time.

As described in Section 3.3.1.2 [Installation directory], page 39, you can select the top installation directory of a software using the GNU build system, when you `./configure` it. All the separate components will be put in their separate sub-directory under that, for example the programs, compiled libraries and library headers will go into `$prefix/bin` (replace `$prefix` with a directory), `$prefix/lib`, and `$prefix/include` respectively. For enhanced modularity, libraries that contain diverse collections of functions (like GSL, WC-SLIB, and Gnuastro), put their header files in a sub-directory unique to themselves. For example all Gnuastro's header files are installed in `$prefix/include/gnuastro`. In your source code, you need to keep the library's sub-directory when including the headers from such libraries, for example `#include <gnuastro/fits.h>`[5]. Not all libraries need to follow this convention, for example CFITSIO only has one header (`fitsio.h`) which is directly installed in `$prefix/include`.

---

[2] The `include/` directory name is taken from the pre-processor's `#include` directive, which is also the motivation behind the 'I' in the `-I` option to the pre-processor.

[3] Try running Gnuastro's `make` and find the directories given to the compiler with the `-I` option.

[4] Nearly every command you see being executed after running `make` is one call to the compiler.

[5] the top `$prefix/include` directory is usually known to the compiler

10.1.2 Linking

To enhance modularity, similar functions are defined in one source file (with a `.c` suffix, see Section 10.1.1 [Headers], page 224, for more). After running `make`, each human-readable, `.c` file is translated (or compiled) into a computer-readable "object" file (ending with `.o`). Note that object files are also created when building programs, they aren't particular to libraries. Try opening Gnuastro's `lib/` and `bin/progname/` directories after running `make` to see these object files[6]. Afterwards, the object files are *linked* together to create an executable program or a library.

The object files contain the full definition of the functions in the respective `.c` file along with a list of any other function (or generally "symbol") that is referenced there. To get a list of those functions you can use the `nm` program which is part of GNU Binutils. For example from the top Gnuastro directory, run:

```
$ nm bin/arithmetic/arithmetic.o
```

This will print a list of all the functions (more generally, 'symbols') that were called within `bin/arithmetic/arithmetic.c` along with some further information (for example a `T` in the second column shows that this function is actually defined here, `U` says that it is undefined here). Try opening the `.c` file to check some of these functions for your self. Run `info nm` for more information.

To recap, the *compiler* created the separate object files mentioned above for each `.c` file. The *linker* will then combine all the symbols of the various object files (and libraries) into one program or library. In the case of Arithmetic (a program) the contents of the object files in `bin/arithmetic/` are copied (and re-ordered) into one final executable file which we can run from the operating system. When the symbols (computer-readable function definitions in most cases) are copied into the output like this, we call the process *static* linking. Let's have a closer look at static linking: we'll assume you have installed Gnuastro into the default `/usr/local/` directory (see Section 3.3.1.2 [Installation directory], page 39). If you tried the `nm` command on one of Arithmetic's object files above, then with the command below you can confirm that all the functions that were defined in the object files (had a `T` in the second column) are also defined in the `astarithmetic` executable:

```
$ nm /usr/local/bin/astarithmetic
```

But you will notice that there are still many undefined symbols (have a `U` in the second column). One class of such functions are Gnuastro's own library functions that start with '`gal_`':

```
$ nm /usr/local/bin/astarithmetic | grep gal_
```

These undefined symbols (functions) will be linked to the executable every time you run arithmetic. Therefore they are known as dynamically *linked* libraries[7]. When the functions of a library need to be dynamically linked, the library is known as a shared library. As we saw above, static linking is done when the executable is being built. However, when a library is linked dynamically, its symbols are only checked with the available libraries at build time: they are not actually copied into the executable. Every time you run the

[6] Gnuastro uses GNU Libtool for portable library creation. Libtool will also make a `.lo` file for each `.c` file when building libraries (`.lo` files are human-readable).

[7] Do not confuse dynamically *linked* libraries with dynamically *loaded* libraries. The former (that is discussed here) are only loaded once at the program startup. However, the latter can be loaded anytime during the program's execution, they are also known as plugins.

program, the linker will be activated and will try to link the program to the installed library before it starts. If you want all the libraries to be statically linked to the executables, you have to tell Libtool (which Gnuastro uses for the linking) to disable shared libraries at configure time[8]:

```
$ configure --disable-shared
```

Try configuring, statically building and installing Gnuastro with the command above. Then check the `gal_` symbols in the installed Arithmetic executable like before. You will see that they are actually copied this time (have a `T` in the second column). If the second column doesn't convince you, look at the executable file size with the following command:

```
$ ls -lh /usr/local/bin/astarithmetic
```

It should be around 4.2 Megabytes with this static linking. If you configure and build Gnuastro again with shared libraries enabled (which is the default), you will notice that it is roughly 100 Kilobytes! This huge difference would have been very significant in the old days, but with the roughly Terabyte storages commonly in use today, it is negligible. Fortunately, output file size is not the only benefit of dynamic linking: since it links to the libraries at run-time (rather than build-time), you don't have to re-build a higher-level program or library when an update comes for one of the lower-level libraries it depends on. You just install the new low-level library and it will automatically be used next time in your higher-level tools. To be fair, this also creates a few complications[9]:

- Reproducibility: Even though your high-level tool has the same version as before, with the updated library, you might not get the same results.
- Broken links: if some functions have been changed or removed in the updated library, then the linker will abort with an error at run-time. Therefore you need to re-build your higher-level program or library.

To see a list of all the shared libraries that are needed for a program or a shared library to run, you can use the GNU C library's `ldd`[10] program, for example:

```
$ ldd /usr/local/bin/astarithmetic
```

Library file names start with a `lib` and end with suffix depending on their type as described below. In between these two is the name of the library, for example `libgnuastro.a` (Gnuastro's static library) and `libgsl.so.0.0.0` (GSL's shared library).

- A static library is known as an archive file and has a `.a` suffix. A static library is not an executable file.
- A shared library ends with a `.so.X.Y.Z` suffix and is executable. The three numbers in the prefix are the version of the shared library. Shared library versions are defined to allow multiple versions of a shared library simultaneously on a system and to help detect possible updates in the library and programs that depend on it by the linker. It is very important to mention that this version number is different from from the software

[8] Libtool is very common and is commonly used. Therefore, you can use this option to configure on most programs using the GNU build system if you want static linking.

[9] Both of these can be avoided by joining the mailing lists of the lower-level libraries and checking the changes in newer versions before installing them. Updates that result in such behaviors are generally heavily emphasized in the release notes.

[10] If your operating system is not using the GNU C library, you might need another tool.

version number (see Section 1.5 [Version numbering], page 5), so do not confuse the two. See the "Library interface versions" chapter of GNU Libtool for more.

For each shared library, we also have two symbolic links ending with `.so.X` and `.so`. They are automatically set by the installer, but you can change them (point them to another version of the library) when you have multiple versions on your system.

For those libraries that use GNU Libtool (including Gnuastro and its dependencies), both static and dynamic libraries are built and installed in the `prefix/lib/` directory (see Section 3.3.1.2 [Installation directory], page 39). In this way other programs can make which ever kind of link that they want.

To link with a library, the linker needs to know where to find the library. You do that with two separate options to the linker (see Section 10.1.3 [Summary and example on libraries], page 229, for an example):

`-L DIR`
: Will tell the linker to look into `DIR` for the libraries. For example `-L/usr/local/lib`, or `-L/home/yourname/.local/lib`. You can make multiple calls to this option, so the linker looks into several directories. Note that the space between `L` and the directory is optional and commonly not used.

`-lLIBRARY`
: Specify the unique name of a library to be linked. As discussed above, library file names have fixed parts which must not be given to this option. So `-lgsl` will guide the linker to either look for `libgsl.a` or `libgsl.so` (depending on the type of linking it is suppose to do). You can link many libraries by repeated calls to this option.

  Very important: The place of this option on the command line matters. This is often a source of confusion for beginners, so let's assume you have asked the linker to link with library A using this option. As soon as the linker confronts this option, it looks into the list of the undefined symbols it has found until that point and does a search in library A for any of those symbols. If any pending undefined symbol is found in library A, it is used. After the search in undefined symbols is complete, the contents of library A are completely discarded from the linker's memory. Therefore, if a later object file or library uses an unlinked symbol in library A, the linker will abort after it has finished its search in all the input libraries or object files.

  As an example, Gnuastro's `gal_array_dlog10_array` function depends on the `log10` function of the C Math library (specified with `-lm`). So the proper way to link something that uses this function is `-lgnuastro -lm`. If instead, you give: `-lm -lgnuastro` the linker will complain and abort.

10.1.3 Summary and example on libraries

After the mostly abstract discussions of Section 10.1.1 [Headers], page 224, and Section 10.1.2 [Linking], page 227, we'll give a small tutorial here. But before that, let's recall the general steps of how your source code is prepared, compiled and linked to the libraries it depends on so you can run it:

1. The **pre-processor** includes the header (`.h`) files into the function definition (`.c`) files, expands pre-processor macros and generally prepares the human-readable source for compilation (reviewed in Section 10.1.1 [Headers], page 224).

2. The **compiler** will translate (compile) the human-readable contents of each source (merged `.c` and the `.h` files, or generally the output of the pre-processor) into the computer-readable code of `.o` files.
3. The **linker** will link the called function definitions from various compiled files to create one unified object. When the unified product has a `main` function, this function is the product's only entry point, enabling the operating system or user to directly interact with it, so the product is a program. When the product doesn't have a `main` function, the linker's product is a library and its exported functions can be linked to other executables (it has many entry points).

The GNU Compiler Collection (or GCC for short) will do all three steps. So as a first example, from Gnuastro's source, go to `tests/lib/`. This directory contains the library tests, you can use these as some simple tutorials. For this demonstration, we will compile and run the `arraymanip.c`. This small program will call Gnuastro library for some simple operations on an array (open it and have a look). To compile this program, run this command inside the directory containing it.

```
$ gcc arraymanip.c -lgnuastro -lm -o arraymanip
```

The two `-lgnuastro` and `-lm` options (in this order) tell GCC to first link with the Gnuastro library and then with C's math library. The `-o` option is used to specify the name of the output executable, without it the output file name will be `a.out` (on most OSs), independent of your input file name(s).

If your top Gnuastro installation directory (let's call it `$prefix`, see Section 3.3.1.2 [Installation directory], page 39) is not recognized by GCC, you will get pre-processor errors for unknown header files. Once you fix it, you will get linker errors for undefined functions. To fix both, you should run GCC as follows: additionally telling it which directories it can find Gnuastro's headers and compiled library (see Section 10.1.1 [Headers], page 224, and Section 10.1.2 [Linking], page 227):

```
$ gcc -I$prefix/include -L$prefix/lib arraymanip.c -lgnuastro -lm     \
      -o arraymanip
```

This single command has done all the pre-processor, compilation and linker operations. Therefore no intermediate files (object files in particular) were created, only a single output executable was created. You are now ready to run the program with:

```
$ ./arraymanip
```

The Gnuastro functions called by this program only needed to be linked with the C math library. But if your program needs WCS coordinate transformations, needs to read a FITS file, needs special math operations (which include its linear algebra operations), or you want it to run on multiple CPU threads, you also need to add these libraries in the call to GCC: `-lgnuastro -lwcs -lcfitsio -lgsl -lgslcblas -pthread -lm`. In Section 10.3 [Gnuastro library], page 233, where each function is documented, it is mentioned which libraries (if any) must also be linked when you call a function. If you feel all these linkings can be confusing, please consider Gnuastro's Section 10.2 [BuildProgram], page 230, program.

10.2 BuildProgram

The number and order of libraries that are necessary for linking a program with Gnuastro library might be too confusing when you need to compile a small program for one particular

job (with one source file). BuildProgram will use the information gathered during configuring Gnuastro and link with all the appropriate libraries on your system. This will allow you to easily compile, link and run programs that use Gnuastro's library with one simple command and not worry about which libraries to link to, or the linking order.

BuildProgram uses GNU Libtool to find the necessary libraries to link against (GNU Libtool is the same program that builds all of Gnuastro's libraries and programs when you run `make`). So in the future, if Gnuastro's prerequisite libraries change or other libraries are added, you don't have to worry, you can just run BuildProgram and internal linking will be done correctly.

10.2.1 Invoking BuildProgram

BuildProgram will compile and link a C source program with Gnuastro's library and all its dependencies, greatly facilitating the compilation and running of small programs that use Gnuastro's library. The executable name is `astbuildprog` with the following general template:

```
$ astbuildprog [OPTION...] C_SOURCE_FILE
```

One line examples:

```
## Compile, link and run 'myprogram.c':
$ astbuildprog myprogram.c

## Similar to previous, but with optimization and compiler warnings:
$ astbuildprog -Wall -O2 myprogram.c

## Compile and link 'myprogram.c', then run it with 'image.fits'
## as its argument:
$ astbuildprog myprogram.c image.fits

## Also look in other directories for headers and linking:
$ astbuildprog -Lother -Iother/dir myprogram.c

## Just build (compile and link) 'myprogram.c', don't run it:
$ astbuildprog --onlybuild myprogram.c
```

If BuildProgram is to run, it needs a C programming language source file as input. By default it will compile and link the program to build the a final executable file and run it. The built executable name can be set with the optional `--output` option. When no output name is set, BuildProgram will use Gnuastro's Section 4.8 [Automatic output], page 77, and remove the suffix of the input and use that as the output name. For the full list of options that BuildProgram shares with other Gnuastro programs, see Section 4.1.2 [Common options], page 52. You may also use Gnuastro's Section 4.2 [Configuration files], page 59, to specify other libraries/headers to use for special directories and not have to type them in every time.

The first argument is considered to be the C source file that must be compiled and linked. Any other arguments (non-option tokens on the command-line) will be passed onto the program when BuildProgram wants to run it. Recall that by default BuildProgram will run the program after building it. This behavior can be disabled with the `--onlybuild` option.

When the `--quiet` option (see Section 4.1.2.3 [Operating mode options], page 55) is not called, BuildPrograms will print the compilation and running commands. Once your program grows and you break it up into multiple files (which are much more easily managed with Make), you can use the linking flags of the non-quiet output in your `Makefile`.

`-I STR`
`--includedir=STR`
: Directory to search for files that you `#include` in your C program. Note that headers relating to Gnuastro and its dependencies don't need this option. This is only necessary if you want to use other headers. It may be called multiple times and order matters. This directory will be searched before those of Gnuastro's build and also the system search directories. See Section 10.1.1 [Headers], page 224, for a thorough introduction.

  From the GNU C Pre-Processor manual: "Add the directory `STR` to the list of directories to be searched for header files. Directories named by `-I` are searched before the standard system include directories. If the directory `STR` is a standard system include directory, the option is ignored to ensure that the default search order for system directories and the special treatment of system headers are not defeated".

`-L STR`
`--linkdir=STR`
: Directory to search for compiled libraries to link the program with. Note that all the directories that Gnuastro was built with will already be used by BuildProgram (GNU Libtool). This option is only necessary if your libraries are in other directories. Multiple calls to this option are possible and order matters. This directory will be searched before those of Gnuastro's build and also the system search directories. See Section 10.1.2 [Linking], page 227, for a thorough introduction.

`-l STR`
`--linklib=STR`
: Library to link with your program. Note that all the libraries that Gnuastro was built with will already be linked by BuildProgram (GNU Libtool). This option is only necessary if you want to link with other directories. Multiple calls to this option are possible and order matters. This library will be linked before Gnuastro's library or its dependencies. See Section 10.1.2 [Linking], page 227, for a thorough introduction.

`-O INT`
`--optimize=INT`
: Compiler optimization level: 0 (for no optimization, good debugging), 1, 2, 3 (for the highest level of optimizations). From the GNU Compiler Collection (GCC) manual: "Without any optimization option, the compiler's goal is to reduce the cost of compilation and to make debugging produce the expected results. Statements are independent: if you stop the program with a breakpoint between statements, you can then assign a new value to any variable or change the program counter to any other statement in the function and get exactly the results you expect from the source code. Turning on optimization flags

makes the compiler attempt to improve the performance and/or code size at the expense of compilation time and possibly the ability to debug the program." See your compiler's manual for the full list of acceptable values to this option.

`-g`
`--debug` Emit extra information in the compiled binary for use by a debugger. When calling this option, it is best to explicitly disable optimization with `-O0`. To combine both options you can run `-gO0` (see Section 4.1.1.2 [Options], page 50, for how short options can be merged into one).

`-W STR`
`--warning=STR`
Print compiler warnings on command-line during compilation. "Warnings are diagnostic messages that report constructions that are not inherently erroneous but that are risky or suggest there may have been an error." (from the GCC manual). It is always recommended to compile your programs with warnings enabled.

All compiler warning options that start with `W` are usable by this option in BuildProgram also, see your compiler's manual for the full list. Some of the most common values to this option are: `pedantic` (Warnings related to standard C), `all` (all issues the compiler confronts).

`-b`
`--onlybuild`
Only build the program, don't run it. By default, the built program is immediately run afterwards.

10.3 Gnuastro library

Gnuastro library's programming constructs (function declarations, macros, data structures, or global variables) are classified by context into multiple header files (see Section 10.1.1 [Headers], page 224)[11]. In this section, the functions in each header will be discussed under a separate sub-section, which includes the name of the header. Assuming a function declaration is in `headername.h`, you can include its declaration in your source code with:

```
# include <gnuastro/headername.h>
```

The names of all constructs in `headername.h` are prefixed with `gal_headername_` (or `GAL_HEADERNAME_` for macros). The `gal_` prefix stands for *G*NU *A*stronomy *L*ibrary.

Gnuastro library functions are compiled into a single file which can be linked on the command-line with the `-lgnuastro` option, (see Section 10.1.2 [Linking], page 227, and Section 10.1.3 [Summary and example on libraries], page 229, for an introduction on linking and example). Gnuastro library is a high-level library which depends on lower level libraries for some operations (see Section 3.1 [Dependencies], page 25). Therefore if at least one of Gnuastro's functions in your program use functions from the dependencies, you will also need to link those dependencies after linking with Gnuastro. The outside libraries that need to be linked for such functions are mentioned following the function name. See Section 10.2 [BuildProgram], page 230, for a small Gnuastro program that will take care of the libraries to link against and lets you focus on your exciting science.

[11] Within Gnuastro's source, all installed `.h` files in `lib/gnuastro/` are accompanied by a `.c` file in `/lib/`.

Libraries are still under heavy development: Gnuastro was initially created to be a collection of command-line programs. However, as the programs and their the shared functions grew, internal (not installed) libraries were added. Since the 0.2 release, the libraries are installable. Hence the libraries are currently under heavy development and will significantly evolve between releases and will become more mature and stable in due time. It will stabilize with the removal of this notice. Check the `NEWS` file for interface changes. If you use the Info version of this manual (see Section 4.7.4 [Info], page 76), you don't have to worry: the documentation will correspond to your installed version.

10.3.1 Configuration information (`config.h`)

The `gnuastro/config.h` header contains information about the full Gnuastro installation on your system. Gnuastro developers should note that this is the only header that is not available within Gnuastro, it is only available to a Gnuastro library user *after* installation. Within Gnuastro, `config.h` (which is included in every Gnuastro `.c` file, see Section 11.3 [Coding conventions], page 323) has more than enough information about the overall Gnuastro installation.

`GAL_CONFIG_VERSION` [Macro]

This macro can be used as a string literal[12] containing the version of Gnuastro that is being used. See Section 1.5 [Version numbering], page 5, for the version formats. For example:

```
printf("Gnuastro version: %s\n", GAL_CONFIG_VERSION);
```

or

```
char *gnuastro_version=GAL_CONFIG_VERSION;
```

`GAL_CONFIG_HAVE_LIBGIT2` [Macro]

Libgit2 is an optional dependency of Gnuastro (see Section 3.1.2 [Optional dependencies], page 28). When it is installed and detected at configure time, this macro will have a value of `1` (one). Otherwise, it will have a value of `0` (zero). Gnuastro also comes with some wrappers to make it easier to use libgit2 (see Section 10.3.22 [Git wrappers (`git.h`)], page 311).

`GAL_CONFIG_HAVE_WCSLIB_VERSION` [Macro]

WCSLIB is the reference library for world coordinate system transformation (see Section 3.1.1.3 [WCSLIB], page 28, and Section 10.3.9 [World Coordinate System (`wcs.h`)], page 277). However, only more recent versions of WCSLIB also provide its version number. If the WCSLIB that is installed on the system provides its version (through the possibly existing `wcslib_version` function), this macro will have a value of one, otherwise it will have a value of zero.

`GAL_CONFIG_HAVE_PTHREAD_BARRIER` [Macro]

The POSIX threads standard define barriers as an optional requirement. Therefore, some operating systems choose to not include it. As one of the `./configure` step checks, Gnuastro we check if your system has this POSIX thread barriers. If so, this macro will have a value of `1`, otherwise it will have a value of `0`. see Section 10.3.2.1 [Implementation of `pthread_barrier`], page 236, for more.

[12] `https://en.wikipedia.org/wiki/String_literal`

`GAL_CONFIG_BIN_OP_UINT8` [Macro]
`GAL_CONFIG_BIN_OP_INT8` [Macro]
`GAL_CONFIG_BIN_OP_UINT16` [Macro]
`GAL_CONFIG_BIN_OP_INT16` [Macro]
`GAL_CONFIG_BIN_OP_UINT32` [Macro]
`GAL_CONFIG_BIN_OP_INT32` [Macro]
`GAL_CONFIG_BIN_OP_UINT64` [Macro]
`GAL_CONFIG_BIN_OP_INT64` [Macro]
`GAL_CONFIG_BIN_OP_FLOAT32` [Macro]
`GAL_CONFIG_BIN_OP_FLOAT64` [Macro]

If biniary arithmetic operators were configured for any type, the respective macro will have a value of `1` (one), otherwise its value will be `0` (zero). Please see the similar configure-time options in Section 3.3.1.1 [Gnuastro configure options], page 37, for a thorough explanation. These are only relevant for you if you intend to use the binary operators of Section 10.3.12 [Arithmetic on datasets (`arithmetic.h`)], page 284,

`GAL_CONFIG_SIZEOF_SIZE_T` [Macro]

The size of (number of bytes in) the system's `size_t` type. Its value is either 4 or 8 for 32-bit and 64-bit systems. You can also get this value with the expression '`sizeof size_t`' without having to include this header.

10.3.2 Multithreaded programming (`threads.h`)

In recent years, newer CPUs don't have significantly higher frequencies any more. However, CPUs are being manufactured with more cores, enabling more than one operation (thread) at each instant. This can be very useful to speed up many aspects of processing and in particular image processing.

Most of the programs in Gnuastro utilize multi-threaded programming for the CPU intensive processing steps. This can potentially lead to a significant decrease in the running time of a program, see Section 4.3.1 [A note on threads], page 62. In terms of reading the code, you don't need to know anything about multi-threaded programming. You can simply follow the case where only one thread is to be used. In these cases, threads are not used and can be completely ignored.

When the C language was defined (the K&R's book was written), using threads was not common, so C's threading capabilities aren't introduced there. Gnuastro uses POSIX threads for multi-threaded programming, defined in the `pthread.h` system wide header. There are various resources for learning to use POSIX threads. An excellent tutorial (`https://computing.llnl.gov/tutorials/pthreads/`) is provided by the Lawrence Livermore National Laboratory, with abundant figures to better understand the concepts, it is a very good start. The book 'Advanced programming in the Unix environment'[13], by Richard Stevens and Stephen Rago, Addison-Wesley, 2013 (Third edition) also has two chapters explaining the POSIX thread constructs which can be very helpful.

An alternative to POSIX threads was OpenMP, but POSIX threads are low level, allowing much more control, while being easier to understand, see Section 11.1 [Why C programming language?], page 320. All the situations where threads are used in Gnuastro

[13] Don't let the title scare you! The two chapters on Multi-threaded programming are very self-sufficient and don't need any more knowledge than K&R.

currently are completely independent with no need of coordination between the threads. Such problems are known as "embarrassingly parallel" problems. They are some of the simplest problems to solve with threads and are also the ones that benefit most from them, see the LLNL introduction[14].

One very useful POSIX thread concept is `pthread_barrier`. Unfortunately, it is only an optional feature in the POSIX standard, so some operating systems don't include it. Therefore in Section 10.3.2.1 [Implementation of `pthread_barrier`], page 236, we introduce our own implementation. This is a rather technical section only necessary for more technical readers and you can safely ignore it. Following that, we describe the helper functions in this header that can greatly simplify writing a multi-threaded program, see Section 10.3.2.2 [Gnuastro's thread related functions], page 237, for more.

10.3.2.1 Implementation of `pthread_barrier`

One optional feature of the POSIX Threads standard is the `pthread_barrier` concept. It is a very useful high-level construct that allows for independent threads to "wait" behind a "barrier" for the rest after they finish. Barriers can thus greatly simplify the code in a multi-threaded program, so they are heavily used in Gnuastro. However, since its an optional feature in the POSIX standard, some operating systems don't include it. So to make Gnuastro portable, we have written our own implementation of those `pthread_barrier` functions.

At `./configure` time, Gnuastro will check if `pthread_barrier` constructs are available on your system or not. If `pthread_barrier` is not available, our internal implementation will be compiled into the Gnuastro library and the definitions and declarations below will be usable in your code with `#include <gnuastro/threads.h>`.

`pthread_barrierattr_t` [Type]

Type to specify the attributes of a POSIX threads barrier.

`pthread_barrier_t` [Type]

Structure defining the POSIX threads barrier.

`int` [Function]
`pthread_barrier_init` (*pthread_barrier_t* `*b`, *pthread_barrierattr_t* `*attr`, *unsigned int* `limit`)

Initialize the barrier `b`, with the attributes `attr` and total `limit` (a number of) threads that must wait behind it. This function must be called before spinning off threads.

`int` [Function]
`pthread_barrier_wait` (*pthread_barrier_t* `*b`)

This function is called within each thread, just before it is ready to return. Once a thread's function hits this, it will "wait" until all the other functions are also finished.

`int` [Function]
`pthread_barrier_destroy` (*pthread_barrier_t* `*b`)

Destroy all the information in the barrier structure. This should be called by the function that spinned-off the threads after all the threads have finished.

[14] `https://computing.llnl.gov/tutorials/parallel_comp/`

Destroy a barrier before re-using it: It is very important to destroy the barrier before (possibly) reusing it. This destroy function not only destroys the internal structures, it also waits (in 1 microsecond intervals, so you will not notice!) until all the threads don't need the barrier structure any more. If you immediately start spinning off new threads with a not-destroyed barrier, then the internal structure of the remaining threads will get mixed with the new ones and you will get very strange and apparently random errors that are extremely hard to debug.

10.3.2.2 Gnuastro's thread related functions

The POSIX Threads functions offered in the C library are very low-level and offer a great range of control over the properties of the threads. So if you are interested in customizing your tools for complicated thread applications, it is strongly encouraged to get a nice familarity with them. Some resources were introduced in Section 10.3.2 [Multithreaded programming (`threads.h`)], page 235.

However, in many cases used in astronomical data analysis, you don't need communication between threads and each target operation can be done independently. Since such operations are very commen, Gnuastro provides the tools below to facilitate the creation and management of jobs without any paritcular knowledge of POSIX Threads for such operations. The most interesting high-level functions of this section are the `gal_threads_number` and `gal_threads_spin_off` that identify the number of threads on the system and spin-off threads. You can see a demonstration of using these functions in Section 10.4.3 [Library demo - multi-threaded operation], page 314.

`gal_threads_params` [C `struct`)]

Structure keeping the parameters of each thread. When each thread is created, a pointer to this structure is passed to it. The `params` element can be the pointer to a structure defined by the user which contains all the necessary parameters to pass onto the worker function. The rest of the elements within this structure are set internally by `gal_threads_spin_off` and are relevant to the worker function.

```
struct gal_threads_params
{
  size_t            id; /* Id of this thread.                  */
  void         *params; /* User-identified pointer.            */
  size_t       *indexs; /* Target indexs given to this thread. */
  pthread_barrier_t *b; /* Barrier for all threads.            */
};
```

`size_t` [Function]
`gal_threads_number ()`

Return the number of threads that the operating system has available for your program. This number is usually fixed for a single machine and doesn't change. So this function is useful when you want to run your program on different machines (with different CPUs).

void [Function]
gal_threads_spin_off (*void* *(*worker)(void *), *void* *caller_params, *size_t* numactions, *size_t* numthreads)

Distribute numactions jobs between numthreads threads and spin-off each thread by calling the worker function. The caller_params pointer will also be passed to worker as part of the gal_threads_params structure. For a fully working example of this function, please see Section 10.4.3 [Library demo - multi-threaded operation], page 314.

void [Function]
gal_threads_attr_barrier_init (*pthread_attr_t* *attr, *pthread_barrier_t* *b, *size_t* limit)

This is a low-level function in case you don't want to use gal_threads_spin_off. It will initialize the general thread attribute attr and the barrier b with limit threads to wait behind the barrier. For maximum efficiency, the threads initialized with this function will be detached. Therefore no communication is possible between these threads and in particular pthread_join won't work on these threads. You have to use the barrier constructs to wait for all threads to finish.

void [Function]
gal_threads_dist_in_threads (*size_t* numactions, *size_t* numthreads, *size_t* **outthrds, *size_t* *outthrdcols)

This is a low-level function in case you don't want to use gal_threads_spin_off. Identify the "index"es (starting from 0) of the actions to be done on each thread in the outthrds array. outthrds is treated as a 2D array with numthreads rows and outthrdcols columns. The indexs in each row, identify the actions that should be done by one thread. Please see the explanation below to understand the purpose of this operation.

Let's assume you have A actions (where there is only one function and the input values differ for each action) and T threads available to the system with $A > T$ (common values for these two would be $A > 1000$ and $T < 10$). Spinning off a thread is not a cheap job and requires a significant number of CPU cycles. Therefore, creating A threads is not the best way to address such a problem. The most efficient way to manage the actions is such that only T threads are created, and each thread works on a list of actions identified for it in series (one after the other). This way your CPU will get all the actions done with minimal overhead.

The purpose of this function is to do what we explained above: each row in the outthrds array contains the indexs of actions which must be done by one thread. outthrds contains outthrdcols columns. In using outthrds, you don't have to know the number of columns. The GAL_BLANK_SIZE_T macro has a role very similar to a string's \0: every row finishes with this macro, so can easily stop parsing the indexes in the row when you confront it. Please see the example program in tests/lib/multithread.c for a demonstration.

10.3.3 Library data types (type.h)

Data in astronomy can have many types, numeric (numbers) and strings (names, identifiers). The former can also be divided into integers and floats, see Section 4.4 [Numeric data types],

page 64, for a thorough discussion of the different numeric data types and which one is useful for different contexts.

To deal with the very large diversity of types that are available (and used in different contexts), in Gnuastro each type is identified with global integer variable with a fixed name, this variable is then passed onto functions that can work on any type or is stored in Gnuastro's Section 10.3.5.1 [Generic data container (`gal_data_t`)], page 245, as one piece of meta-data.

The actual values within these integer constants is irrelevant and you should never rely on them. When you need to check, explicitly use the named variable in the table below. If you want to check with more than one type, you can use C's `switch` statement.

Since Gnuastro heavily deals with file input-output, the types it defines are fixed width types, these types are portable to all systems and are defined in the standard C header `stdint.h`. You don't need to include this header, it is included by any Gnuastro header that deals with the different types. However, the most commonly used types in a C (or C++) program (for example `int` or `long` are not defined by their exact width (storage), but by their minimum storage. So for example on some systems, `int` may be 2 bytes (16-bits, the minimum required by the standard) and on others it may be 4 bytes (32-bits, common in modern systems).

With every type, a unique "blank" value (or place holder showing the absence of data) can be defined. Please see Section 10.3.4 [Library blank values (`blank.h`)], page 243, for constants that Gnuastro recognizes as a blank value for each type. See Section 4.4 [Numeric data types], page 64, for more explanation on the limits and particular aspects of each type.

`GAL_TYPE_INVALID` [Global integer]
: This is just a place holder to specifically mark that no type has been set.

`GAL_TYPE_BIT` [Global integer]
: Identifier for a bit-stream. Currently no program in Gnuastro works directly on bits, but features will be added in the future.

`GAL_TYPE_UINT8` [Global integer]
: Identifier for an unsigned, 8-bit integer type: `uint8_t` (from `stdint.h`), or an `unsigned char` in most modern systems.

`GAL_TYPE_INT8` [Global integer]
: Identifier for a signed, 8-bit integer type: `int8_t` (from `stdint.h`), or an `signed char` in most modern systems.

`GAL_TYPE_UINT16` [Global integer]
: Identifier for an unsigned, 16-bit integer type: `uint16_t` (from `stdint.h`), or an `unsigned short` in most modern systems.

`GAL_TYPE_INT16` [Global integer]
: Identifier for a signed, 16-bit integer type: `int16_t` (from `stdint.h`), or a `short` in most modern systems.

`GAL_TYPE_UINT32` [Global integer]
: Identifier for an unsigned, 32-bit integer type: `uint32_t` (from `stdint.h`), or an `unsigned int` in most modern systems.

`GAL_TYPE_INT32` [Global integer]
Identifier for a signed, 32-bit integer type: `int32_t` (from `stdint.h`), or an `int` in most modern systems.

`GAL_TYPE_UINT64` [Global integer]
Identifier for an unsigned, 64-bit integer type: `uint64_t` (from `stdint.h`), or an `unsigned long` in most modern 64-bit systems.

`GAL_TYPE_INT64` [Global integer]
Identifier for a signed, 64-bit integer type: `int64_t` (from `stdint.h`), or an `long` in most modern 64-bit systems.

`GAL_TYPE_SIZE_T` [Global integer]
Identifier for a `size_t` type. This is just an alias to `uint32`, or `uint64` types for 32-bit, or 64-bit systems respectively.

`GAL_TYPE_FLOAT32` [Global integer]
Identifier for a 32-bit single precision floating point type or `float` in C.

`GAL_TYPE_FLOAT64` [Global integer]
Identifier for a 64-bit double precision floating point type or `double` in C.

`GAL_TYPE_COMPLEX32` [Global integer]
Identifier for a complex number composed of two `float` types. Note that the complex type is not yet fully implemented in all Gnuastro's programs.

`GAL_TYPE_COMPLEX64` [Global integer]
Identifier for a complex number composed of two `double` types. Note that the complex type is not yet fully implemented in all Gnuastro's programs.

`GAL_TYPE_STRING` [Global integer]
Identifier for a string of characters (`char *`).

`GAL_TYPE_STRLL` [Global integer]
Identifier for a linked list of string of characters (`gal_list_str_t`, see Section 10.3.7.1 [List of strings], page 257).

The functions below are defined to make working with the integer constants above easier. In the functions below, the constants above can be used for the `type` input argument.

`size_t` [Function]
`gal_type_sizeof` (*uint8_t* `type`)
Return the number of bytes occurpied by `type`. Internally, this function uses C's `sizeof` operator to measure the size of each type.

`char *` [Function]
`gal_type_name` (*uint8_t* `type`, *int* `long_name`)
Return a string that contains the name of `type`. This can be used in messages to the users when your function/program accepts many types. It can return both short and long formats of the type names (for example `f32` and `float32`). If `long_name` is non-zero, the long format will be returned, otherwise the short name will be returned.

The output string is statically allocated, so it should not be freed. This function is the inverse of the `gal_type_from_name` function. For the full list of names/strings that this function will return, see Section 4.4 [Numeric data types], page 64.

`uint8_t` [Function]
`gal_type_from_name` (*char* `*str`)

Return the Gnuastro integer constant that corresponds to the string `str`. This function is the inverse of the `gal_type_name` function and accepts both the short and long formats of each type. For the full list of names/strings that this function will return, see Section 4.4 [Numeric data types], page 64.

`void` [Function]
`gal_type_min` (*uint8_t* `type`, *void* `*in`)

Put the minimum possible value of `type` in the space pointed to by `in`. Since the value can have any type, this function doesn't return anything, it assumes the space for the given type is available to `in` and writes the value there. Here is one example

```
int32_t min;
gal_type_min(GAL_TYPE_INT32, &min);
```

Note: Do not use the minimum value for a blank value of a general (initially unknown) type, please use the constants/functions provided in Section 10.3.4 [Library blank values (`blank.h`)], page 243, for the definition and usage of blank values.

`void` [Function]
`gal_type_max` (*uint8_t* `type`, *void* `*in`)

Put the maximum possible value of `type` in the space pointed to by `in`. Since the value can have any type, this function doesn't return anything, it assumes the space for the given type is available to `in` and writes the value there. Here is one example

```
uint16_t max;
gal_type_max(GAL_TYPE_INT16, &max);
```

Note: Do not use the maximum value for a blank value of a general (initially unknown) type, please use the constants/functions provided in Section 10.3.4 [Library blank values (`blank.h`)], page 243, for the definition and usage of blank values.

`int` [Function]
`gal_type_is_list` (*uint8_t* `type`)

Return 1 if the type is a linked list and zero otherwise.

`int` [Function]
`gal_type_out` (*int* `first_type`, *int* `second_type`)

Return the larger of the two given types which can be used for the type of the output of an operation involving the two input types.

`char *` [Function]
`gal_type_bit_string` (*void* `*in`, *size_t* `size`)

Return the bit-string in the `size` bytes that `in` points to. The string is dynamically allocated and must be freed afterwards. You can use it to inspect the bits within one region of memory. Here is one short example:

```
int32_t a=2017;
```

```
char *bitstr=gal_type_bit_string(&a, 4);
printf("%d: %s (%X)\n", a, bitstr, a);
free(bitstr);
```

which will produce:

```
2017: 11100001000001110000000000000000  (7E1)
```

As the example above shows, the bit-string is not the most efficient way to inspect bits. If you are familiar with hexademical notation, it is much more compact, see `https://en.wikipedia.org/wiki/Hexadecimal`. You can use `printf`'s `%x` or `%X` to print integers in hexadecimal format.

`char *` [Function]
`gal_type_to_string` (*void* `*ptr`, *uint8_t* `type`, *int* `quote_if_str_has_space`);

Read the contents of the memory that `ptr` points to (assuming it has type `type` and print it into an allocated string which is returned.

If the memory is a string of characters and `quote_if_str_has_space` is non-zero, the output string will have double-quotes around it if it contains space characters. Also, note that in this case, `ptr` must be a pointer to an array of characters (or `char **`), as in the example below (which will put `"sample string"` into `out`):

```
char *out, *string="sample string"
out = gal_type_to_string(&string, GAL_TYPE_STRING, 1);
```

`int` [Function]
`gal_type_from_string` (*void* `**out`, *char* `*string`, *uint8_t* `type`)

Read a string as a given data type and put a the pointer to it in `*out`. When `*out!=NULL`, then it is assumed to be already allocated and the value will be simply put the memory. If `*out==NULL`, then space will be allocated for the given type and the string will be read into that type.

Note that when we are dealing with a string type, `*out` should be interpretted as `char **` (one element in an array of pointers to different strings). In other words, `out` should be `char ***`.

This function can be used to fill in arrays of numbers from strings (in an already allocated data structure), or add nodes to a linked list (if the type is a list type). For an array, you have to pass the pointer to the `i`th element where you want the value to be stored, for example `&(array[i])`.

If the string was successfully parsed to the requested type, this function will return a `0` (zero), otherwise it will return `1` (one). This output format will help you check the status of the conversion in a code like the example below:

```
if( gal_type_from_string(&out, string, GAL_TYPE_FLOAT32) )
  {
    fprintf(stderr, "%s couldn't be read as float32.\n", string);
    exit(EXIT_FAILURE);
  }
```

`void *` [Function]
`gal_type_string_to_number` (*char* `*string`, *uint8_t* `*type`)

Read `string` into smallest type that can host the number, the allocated space for the number will be returned and the type of the number will be put into the memory

that `type` points to. If `string` couldn't be read as a number, this function will return `NULL`.

For the ranges acceptable by each type see Section 4.4 [Numeric data types], page 64. For integers it is clear, for floating point types, this function will count the number of significant digits and determine if the given string is single or double precision as described in that section.

10.3.4 Library blank values (`blank.h`)

When the position of an element in a dataset is important (for example a pixel in an image), a place-holder is necessary for the element if we don't have a value to fill it with (for example the CCD cannot read those pixels). We cannot simply shift all the other pixels to fill in the one we have no value for. In other cases, it often occurs that the field of sky that you are studying is not a clean rectangle to nicely fit into the boundaries of an image. You need a way to separate the pixels outside your scientific field from those inside it. Blank values act as these place holders in a dataset. They have no usable value but they have a position.

Every type needs a corresponding blank value (see Section 4.4 [Numeric data types], page 64, and Section 10.3.3 [Library data types (`type.h`)], page 238). Floating point types have a unique value identified by IEEE known as Not-a-Number (or NaN) which is a unique value that is recognized by the compiler. However, integer and string types don't have any standard value. For integers, in Gnuastro we take an extremum of the given type: for signed types (that allow negatives), the minimum possible value is used as blank and for unsigned types (that only accept positives), the maximum possible value is used. To be generic and easy to read/write we define a macro for these blank values and strongly encourage you only use these, and never make any assumption on the value of a type's blank value.

The IEEE NaN blank value type is defined to fail on any comparison, so if you are dealing with floating point types, you cannot use equality (a NaN will *not* be equal to a NaN). If you know your dataset if floating point, you can use the `isnan` function in C's `math.h` header. For a description of numeric data types see Section 4.4 [Numeric data types], page 64. For the constants identifying integers, please see Section 10.3.3 [Library data types (`type.h`)], page 238.

`GAL_BLANK_UINT8` [Global integer]
: Blank value for an unsigned, 8-bit integer.

`GAL_BLANK_INT8` [Global integer]
: Blank value for a signed, 8-bit integer.

`GAL_BLANK_UINT16` [Global integer]
: Blank value for an unsigned, 16-bit integer.

`GAL_BLANK_INT16` [Global integer]
: Blank value for a signed, 16-bit integer.

`GAL_BLANK_UINT32` [Global integer]
: Blank value for an unsigned, 32-bit integer.

`GAL_BLANK_INT32` [Global integer]
: Blank value for a signed, 32-bit integer.

`GAL_BLANK_UINT64` [Global integer]
Blank value for an unsigned, 64-bit integer.

`GAL_BLANK_INT64` [Global integer]
Blank value for a signed, 64-bit integer.

`GAL_BLANK_SIZE_T` [Global integer]
Blank value for `size_t` type (`uint32_t` or `uint64_t` in 32-bit or 64-bit systems).

`GAL_BLANK_FLOAT32` [Global integer]
Blank value for a single precision, 32-bit floating point type (IEEE NaN value).

`GAL_BLANK_FLOAT64` [Global integer]
Blank value for a double precision, 64-bit floating point type (IEEE NaN value).

`GAL_BLANK_STRING` [Global integer]
Blank value for string types (this is itsself a string, it isn't the `NULL` pointer).

The functions below can be used to work with blank pixels.

`void` [Function]
`gal_blank_write` (*void* `*pointer`, *uint8_t* `type`)
Write the blank value for the given `type` into the space that `pointer` points to. This can be used when the space is already allocated (for example one element in an array or a statically allocated variable).

`void *` [Function]
`gal_blank_alloc_write` (*uint8_t* `type`)
Allocate the space required to keep the blank for the given data type `type`, write the blank value into it and return the pointer to it.

`void` [Function]
`gal_blank_initialize` (*gal_data_t* `*input`)
Initialize all the elements in the `input` dataset to the blank value that corresponds to its type. If `input` is a tile over a larger dataset, only the region that the tile covers will be set to blank.

`int` [Function]
`gal_blank_present` (*gal_data_t* `*input`, *int* `updateflag`)
Return 1 if the dataset has a blank value and zero if it doesn't. Before checking the dataset, this function will look at `input`'s flags. If the `GAL_DATA_FLAG_HASBLANK` or `GAL_DATA_FLAG_DONT_CHECK_ZERO` bits of `input->flag` are set to 1, this function will not do any check and will just use the information in the flags. This can greatly speed up processing when a dataset needs to be checked multiple times.

If you want to re-check a dataset which has non-zero flags, then explicitly set the appropriate flag to zero before calling this function. When there are no other flags, you can just set `input->flags` to zero, otherwise you can use this expression:

```
input->flags &= ~(GAL_DATA_FLAG_HASBLANK | GAL_DATA_FLAG_USE_ZERO);
```

When `updateflags` is zero, this function has no side-effects on the dataset: it will not toggle the flags. When the dataset's flags were not used and `updateflags` is non-zero, this function will set the flags appropriately to avoid having to re-check the dataset in future calls.

`gal_data_t *` [Function]
`gal_blank_flag` (*gal_data_t* `*input`)

Create a dataset of the the same size as the input, but with an `uint8_t` type that has a value of 1 for data that are blank and 0 for those that aren't.

`void` [Function]
`gal_blank_remove` (*gal_data_t* `*input`)

Remove blank elements from a dataset, convert it to a 1D dataset, and adjust the size properly (the number of non-blank elements). In practice this function doesn't `realloc` the input array, it just shifts the blank elements to the end and adjusts the size elements of the `gal_data_t`, see Section 10.3.5.1 [Generic data container (`gal_data_t`)], page 245.

`char *` [Function]
`gal_blank_as_string` (*uint8_t* `type`, *int* `width`)

Write the blank value for the given data type `type` into a string and return it. The space for the string is dynamically allocated so it must be freed after you are done with it.

10.3.5 Data container (`data.h`)

Astronomical datasets have various dimensions, for example 1D spectra or table columns, 2D images, or 3D Integral field data cubes. Datasets can also have various numeric data types, depending on the operation/purpose, for example processed images are in commonly in floating point, but their mask images are integers using bit-wise flags to identify certain classes of special pixels, see Section 4.4 [Numeric data types], page 64). Certain other information about a dataset are also commonly necessary, for example the units of the dataset, the name of the dataset and some comments. To deal with any generic dataset, Gnuastro defines the `gal_data_t` as input or output.

10.3.5.1 Generic data container (`gal_data_t`)

To be able to deal with any dataset (various dimensions, numeric data types, units and higher-level structures), Gnuastro defines the `gal_data_t` type which is the input/output container of choice for many of Gnuastro library's functions. It is defined in `gnuastro/data.h`. If you will be using ('`# include`'ing) those libraries, you don't need to include this header explicity, it is already included by any library header that uses `gal_data_t`.

`gal_data_t` [Type (C `struct`)]

The main container for datasets in Gnuastro. It can host data of any dimensionality, with any numeric data type. It is actually a structure, but `typedef`'d as a new type to avoid having to write the `struct` before any declaration. The actual structure is shown below which is followed by a description of each element.

```
typedef struct gal_data_t
{
  void     *restrict array;  /* Basic array information.   */
  uint8_t             type;
  size_t              ndim;
```

```
  size_t            *dsize;
  size_t              size;
  char           *mmapname;
  size_t        minmapsize;

  int                 nwcs;  /* WCS information.           */
  struct wcsprm       *wcs;

  uint8_t             flag;  /* Content description.       */
  int               status;
  char               *name;
  char               *unit;
  char            *comment;

  int             disp_fmt;  /* For text printing.         */
  int           disp_width;
  int       disp_precision;

  struct gal_data_t  *next;  /* For higher-level datasets. */
  struct gal_data_t *block;
} gal_data_t;
```

The list below contains a description for each `gal_data_t` element.

`void *restrict array`

This is the pointer to the main array of the dataset containing the raw data (values). All the other elements in this data-structure are actually meta-data enabling us to use/understand the series of values in this array. It must allow data of any type (see Section 4.4 [Numeric data types], page 64), so it is defined as a `void *` pointer. A `void *` array is not directly usable in C, so you have to cast it to proper type before using it, please see Section 10.4.1 [Library demo - reading a FITS image], page 312, for a demonstration.

The `restrict` keyword was formally introduced in C99 and is used to tell the compiler that at any moment only this pointer will modify what it points to (the a pixel in an image for example)[15]. This extra piece of information can greatly help in compiler optimizations and thus the running time of the program. But older compilers might not have this capability, so at `./configure` time, Gnuastro checks this feature and if the user's compiler doesn't support `restrict`, it will be removed from this definition.

`uint8_t type`

A fixed code (integer) used to identify the type of data in `array` (see Section 4.4 [Numeric data types], page 64). For the list of acceptable values to this variable, please see Section 10.3.3 [Library data types (`type.h`)], page 238.

`size_t ndim`

The dataset's number of dimensions.

[15] Also see `https://en.wikipedia.org/wiki/Restrict`.

`size_t *dsize`

The size of the dataset along each dimension. This is an array (with `ndim` elements), of positive integers in row-major order[16] (based on C). When a data file is read into memory with Gnuastro's libraries, this array is dynamically allocated based on the number of dimensions that the dataset has.

It is important to remember that C's row-major ordering is the opposite of the FITS standard which is in column-major order: in the FITS standard the fastest dimension's size is specified by `NAXIS1`, and slower dimensions follow. The FITS standard was defined mainly based on the Fortran language which is the opposite of C's approach to multi-dimensional arrays (and also starts counting from 1 not 0). Hence if a FITS image has `NAXIS1==20` and `NAXIS2==50`, the `dsize` array must be filled with `dsize[0]==50` and `dsize[1]==20`.

The fastest dimension is the one that is contiguous in memory: to increment by one along that dimension, just go to the next element in the array. As we go to slower dimensions, the number of memory cells we have to skip for an increment along that dimension becomes larger.

`size_t size`

The total number of elements in the dataset. This is actually a multiplication of all the values in the `dsize` array, so it is not an indepenent parameter. However, low-level operations with the dataset (irrespective of its dimensionality) commonly need this number, so this element is designed to avoid calculating it everytime.

`char *mmapname`

Name of file hosting the `mmap`'d contents of `array`. If the value of this variable is `NULL`, then the contents of `array` are actually stored in RAM, not in a file on the HDD/SSD. See the description of `minmapsize` below for more.

If a file is used, it will be kept in the hidden `.gnuastro` directory with a randomly selected name to allow multiple arrays to be kept there at the same time. When `gal_data_free` is called the randomly named file will be deleted.

`size_t minmapsize`

The minimum size of an array (in bytes) to store the contents of `array` as a file (on the non-volatile HDD/SSD), not in RAM. This can be very useful for large datasets which can be very memory intensive and the user's hardware RAM might not be sufficent to keep/process it. A random filename is assigned to the array which is available in the `mmapname` element of `gal_data_t` (above), see there for more.

When this variable has a value of `0` (zero), any allocated `array` will actually be in a file (not in RAM). When the value is `-1` (largest possible number in the unsigned types including `size_t`) the array will be definitely allocated in RAM.

Please note that using a non-volatile file instead of RAM will significantly increase the programs running time, especially on HDDs. So it is best to give this option very large values (depending on how much memory you will need

[16] Also see `https://en.wikipedia.org/wiki/Row-_and_column-major_order`.

for a given input). For example your processing might involve a copy of the the input (possibly to a wider data type which takes more bytes for each element), so take all such issues into consideration. `minmapsize` is actually stored in each `gal_data_t`, so it can be passed on to derivate datasets.

`nwcs` The number of WCS coordinate representations (for WCSLIB).

`struct wcsprm *wcs`
: The main WCSLIB structure keeping all the relevant information necessary for WCSLIB to do its processing and convert data-set positions into real-world positions. When it is given a `NULL` value, all posssible WCS calculations/measurements will be ignored.

`uint8_t flag`
: Bit-wise flags to describe general properties of the dataset. The number of bytes available in this flag is stored in the `GAL_DATA_FLAG_SIZE` macro. Note that you should use bit-wise operators[17] to check these flags. The currently recognized bits are stored in these macros:

  `GAL_DATA_FLAG_BLANK_CH`
  : Marking that the dataset has been checked for blank values. Therefore, the value of the bit in `GAL_DATA_FLAG_HASBLANK` is relibable. Without this bit, when a dataset doesn't have any blank values (and this has been checked), the `GAL_DATA_FLAG_HASBLANK` bit will be zero so a checker has no way to know if this zero is real or if no check has been done yet.

  `GAL_DATA_FLAG_HASBLANK`
  : This bit has a value of `1` when the given dataset has blank values. If this bit is `0` and `GAL_DATA_FLAG_BLANK_CH` is `1`, then the dataset has been checked and it didn't have any blank values, so there is no more need for further checks.

  `GAL_DATA_FLAG_SORT_CH`
  : Marking that the dataset is already checked for being sorted or not and thus that the possible `0` values in `GAL_DATA_FLAG_SORTED_I` and `GAL_DATA_FLAG_SORTED_D` are meaningful.

  `GAL_DATA_FLAG_SORTED_I`
  : This bit has a value of `1` when the given dataset is sorted in an increasing manner. If this bit is `0` and `GAL_DATA_FLAG_SORT_CH` is `1`, then the dataset has been checked and wasn't sorted (increasing), so there is no more need for further checks.

  `GAL_DATA_FLAG_SORTED_D`
  : This bit has a value of `1` when the given dataset is sorted in a decreasing manner. If this bit is `0` and `GAL_DATA_FLAG_SORT_CH` is `1`, then the dataset has been checked and wasn't sorted (decreasing), so there is no more need for further checks.

[17] See `https://en.wikipedia.org/wiki/Bitwise_operations_in_C`.

The macro `GAL_DATA_FLAG_MAXFLAG` contains the largest internally used bit-position. Higher-level flags can be defined with the bit-wise shift operators using this macro to define internal flags for libraries/programs that depend on Gnuastro without causing any possible conflict with the internal flags discussed above or having to check the values manually on every release.

`int status`
: A context-specific status values for this data-structure. This integer will not be set by Gnuastro's libraries. You can use it keep some additional information about the dataset (with integer constants) depending on your applications.

`char *name`
: The name of the dataset. If the dataset is a multi-dimensional array and read/written as a FITS image, this will be the value in the `EXTNAME` FITS keyword. If the dataset is a one-dimensional table column, this will be the column name. If it is set to `NULL` (by default), it will be ignored.

`char *unit`
: The units of the dataset (for example `BUNIT` in the standard FITS keywords) that will be read from or written to files/tables along with the dataset. If it is set to `NULL` (by default), it will be ignored.

`char *comment`
: Any further explanation about the dataset which will be written to any output file if present.

`disp_fmt`
: Format to use for printing each element of the dataset to a plain text file, the acceptable values to this element are defined in Section 10.3.11 [Table input output (`table.h`)], page 281. Based on C's `printf` standards.

`disp_width`
: Width of printing each element of the dataset to a plain text file, the acceptable values to this element are defined in Section 10.3.11 [Table input output (`table.h`)], page 281. Based on C's `printf` standards.

`disp_precision`
: Width of printing each element of the dataset to a plain text file, the acceptable values to this element are defined in Section 10.3.11 [Table input output (`table.h`)], page 281. Based on C's `printf` standards.

`gal_data_t *next`
: Through this pointer, you can link a `gal_data_t` with other datasets related datasets, for example the different columns in a dataset each have one `gal_data_t` associate with them and they are linked to each other using this element. There are several functions described below to facilitate using `gal_data_t` as a linked list. See Section 10.3.7 [Linked lists (`list.h`)], page 255, for more on these wonderful high-level constructs.

`gal_data_t *block`
: Pointer to the start of the complete allocated block of memory. When this pointer is not `NULL`, the dataset is not treated as a contiguous patch of memory. Rather, it is seen as covering only a portion of the larger patch of memory that

`block` points to. See Section 10.3.13 [Tessellation library (`tile.h`)], page 288, for a more thorough explanation and functions to help work with tiles that are created from this pointer.

10.3.5.2 Dataset size and allocation

Gnuastro's main data container was defined in Section 10.3.5.1 [Generic data container (`gal_data_t`)], page 245. The functions listed in this section describe the most basic operations on `gal_data_t`: those related to the size, pointers, allocation and freeing. These functions are declared in `gnuastro/data.h` which is also visible from the function names (see Section 10.3 [Gnuastro library], page 233).

`int` [Function]
`gal_data_dsize_is_different` (*gal_data_t* `*first`, *gal_data_t* `*second`)

Return `1` (one) if the two datasets don't have the same size along all dimensions. This function will also return `1` when the number of dimensions of the two datasets are different.

`void *` [Function]
`gal_data_ptr_increment` (*void* `*pointer`, *size_t* `increment`, *uint8_t* `type`)

Return a pointer to an element that is `increment` elements ahead of `pointer`, assuming each element has type of `type` (for the type codes, see Section 10.3.3 [Library data types (`type.h`)], page 238).

When working with the `array` elements of `gal_data_t`, we are actually dealing with `void *` pointers. However, pointer arithmetic doesn't apply to `void *`, because the system doesn't know how many bytes there are in each element to increment the pointer respectively. This function will use the given `type` to calculate where the incremented element is located in memory.

`size_t` [Function]
`gal_data_ptr_dist` (*void* `*earlier`, *void* `*later`, *uint8_t* `type`)

Return the number of elements between `earlier` and `later` assuming each element has a type defined by `type` (for the type codes, see Section 10.3.3 [Library data types (`type.h`)], page 238).

`void *` [Function]
`gal_data_malloc_array` (*uint8_t* `type`, *size_t* `size`, *const char* `*funcname`, *const char* `*varname`)

Allocate an array of type `type` with `size` elements in RAM (for the type codes, see Section 10.3.3 [Library data types (`type.h`)], page 238). This is effectively just a wrapper around C's `malloc` function but takes Gnuastro's integer type codes and will also abort with an error if there the allocation was not successful.

When space cannot be allocated, this function will abort the program with a message containing the reason for the failure. `funcname` (name of the function calling this function) and `varname` (name of variable that needs this space) will be used in this error message if they are not `NULL`. In most modern compilers, you can use the generic `__func__` variable for `funcname`. In this way, you don't have to manually copy and paste the function name or worry about it changing later (`__func__` was standardized in C99).

void * [Function]
gal_data_calloc_array (*uint8_t* type, *size_t* size), *const char* *funcname, *const char* *varname)

Similar to gal_data_malloc_array, but the space is cleared (set to 0) after allocation.

void [Function]
gal_data_initialize (*gal_data_t* *data, *void* *array, *uint8_t* type, *size_t* ndim, *size_t* *dsize, *struct wcsprm* *wcs, *int* clear, *size_t* minmapsize, *char* *name, *char* *unit, *char* *comment)

Initialize the given data structure (data) with all the given values. Note that the raw input gal_data_t must already have been allocated before calling this function. For a description of each variable see Section 10.3.5.1 [Generic data container (gal_data_t)], page 245. It will set the values and do the necessary allocations. If they aren't NULL, all input arrays (dsize, wcs, name, unit, comment) are separately copied (allocated) by this function for usage in data, so you can safely use one value to initialize many datasets or use statically allocated variables in this function call. Once you are done with the dataset, you can clean all the allocated spaces with gal_data_free_contents.

If array is not NULL, it will be directly copied into data->array and no new space will be allocated for the array of this dataset, this has many low-level advantages and can be used to work on regions of a dataset instead of the whole allocated array (see the description under block in Section 10.3.5.1 [Generic data container (gal_data_t)], page 245, for one example). If the given pointer is not the start of an allocated block of memory or it is used in multiple datasets, be sure to set it to NULL (with data->array=NULL) before cleaning up with gal_data_free_contents.

void * [Function]
gal_data_alloc (*void* *array, *uint8_t* type, *size_t* ndim, *size_t* *dsize, *struct wcsprm* *wcs, *int* clear, *size_t* minmapsize, *char* *name, *char* *unit, *char* *comment)

Dynamically allocate a gal_data_t and initialize it will all the given values. See the description of gal_data_initialize and Section 10.3.5.1 [Generic data container (gal_data_t)], page 245, for more information. This function will often be the most frequently used because it allocates the gal_data_t hosting all the values *and* initializes it. Once you are done with the dataset, be sure to clean up all the allocated spaces with gal_data_free.

void [Function]
gal_data_free_contents (*gal_data_t* *data)

Free all the non-NULL pointers in gal_data_t, for a complete description of the gal_data_t contents, see Section 10.3.5.1 [Generic data container (gal_data_t)], page 245.

void [Function]
gal_data_free (*gal_data_t* *data)

Free all the non-NULL pointers in gal_data_t, then free the actual data structure.

10.3.5.3 Arrays of datasets

Gnuastro's generic data container (`gal_data_t`) is a very versatile structure that can be used in many higher-level contexts. One such higher-level construct is an array of `gal_data_t` structures to simplify the allocation (and later clearning) of several `gal_data_ts` that are related.

For example, each column in a table is usually represented by one `gal_data_t` (so it has its own name, numeric data type, units and etc). A table (with many columns) can be seen as an array of `gal_data_ts` (when the number of columns is known a-priori). The functions below are defined to create a cleared array of data structures and to free them in the end. These functions are declared in `gnuastro/data.h` which is also visible from the function names (see Section 10.3 [Gnuastro library], page 233).

`gal_data_t *` [Function]
`gal_data_array_calloc` (*size_t* `size`)

Allocate an array of `gal_data_t` with `size` elements. This function will also initialize all the values (`NULL` for pointers and 0 for other types). You can use `gal_data_initialize` to fill each element of the array afterwards. The following code snippet is one example of doing this.

```
size_t i;
gal_data_t *dataarr;
dataarr=gal_data_array_calloc(10);
for(i=0;i<10;++i) gal_data_initialize(&dataarr[i], ...);
...
gal_data_array_free(dataarr, 10, 1);
```

`void` [Function]
`gal_data_array_free` (*gal_data_t* `*dataarr`, *size_t* `num`, *int* `free_array`)

Free all the `num` elements within `dataarr` and the actual allocated array. If `free_array` is not zero, then the `array` element of all the datasets will also be freed, see Section 10.3.5.1 [Generic data container (`gal_data_t`)], page 245.

10.3.5.4 Copying datasets

The functions in this section describes Gnuastro's facilities to copy a given datset into another. The new dataset can have a different type (including a string), it can be already allocated (in which case only the values will be written into it). In all these cases, if the input dataset is a tile, only the data within the tile are copied.

In many of the functions here, it is possible to copy the dataset to a new numeric data type (see Section 4.4 [Numeric data types], page 64. In such cases, Gnuastro's library is going to use the native conversion by C. So if you are converting to a smaller type, it is up to you to make sure that the values fit into the output type.

`gal_data_t *` [Function]
`gal_data_copy` (*gal_data_t* `*in`)

Return a new dataset that is a copy of `in`, the main meta-data of the input is also copied into the output.

gal_data_t * [Function]
gal_data_copy_to_new_type (*gal_data_t* *in, *uint8_t* newtype)

Return a copy of the dataset `in`, converted to `newtype`, see Section 10.3.3 [Library data types (type.h)], page 238, for Gnuastro library's type identifiers. The returned dataset will have all meta-data accept their type equal to the input's metadata.

gal_data_t * [Function]
gal_data_copy_to_new_type_free (*gal_data_t* *in, *uint8_t* newtype)

Return a copy of the dataset `in` that is converted to `newtype` and free the input dataset. See Section 10.3.3 [Library data types (type.h)], page 238, for Gnuastro library's type identifiers. The returned dataset will have all meta-data, except their type, equal to the input's metadata. This function is similar to `gal_data_copy_to_new_type`, except that it will free the input dataset.

void [Function]
gal_data_copy_to_allocated (*gal_data_t* *in, *gal_data_t* *out)

Copy the contents of the array in `in` into the already allocated array in `out`. The types of the input and output may be different, type conversion will be done internally. When `in->size != out->size` this function will behave as follows:

`out->size < in->size`
: This function won't re-allocate the necessary space, it will abort with an error, so please check before calling this function.

`out->size > in->size`
: This function will write the values in `out->size` and `out->dsize` from the same values of `in`. So if you want to use a pre-allocated space/dataset multiple times with varying input sizes, be sure to reset `out->size` before every call to this function.

gal_data_t * [Function]
gal_data_copy_string_to_number (*char* *string)

Read `string` into the smallest type that can store the value (see Section 4.4 [Numeric data types], page 64). This function is just a wrapper for the `gal_type_string_to_number`, but will put the value into a single-element dataset.

10.3.6 Dimensions (dimension.h)

An array is a contiguous region of memory. Hence, at the lowest level, every element of an array just has one single-valued position: the number of elements that lie between it and the first element in the array. This is also known as the *index* of the element within the array. Dimensionality is high-level abstraction (meta-data) that we project onto that contiguous patch of memory. When the array is interpreted as a one-dimensional dataset, this index is also the *coordinate* of the element. But once we associate the patch of memory with a higher dimensionality, there must also be one coordinate for each dimension.

The functions and macros in this section provide you with the tools to convert an index into a coordinate and vice-versa along with several other issues for example issues with the neighbors of an element in a multi-dimensional context.

size_t [Function]
gal_dimension_total_size (*size_t* ndim, *size_t* *dsize)

Return the total number of elements for a dataset with ndim dimensions that has dsize elements along each dimension.

size_t * [Function]
gal_dimension_increment (*size_t* ndim, *size_t* *dsize)

Return an allocated array that has the number of elements necessary to increment an index along every dimension. For example along the fastest dimension (last element in the dsize and returned arrays), the value is 1 (one).

size_t [Function]
gal_dimension_num_neighbors (*size_t* ndim)

The maximum number of neighbors (any connectivity) that a data element can have in ndim dimensions. Effectively, this function just returns $3^n - 1$ (where n is the number of dimensions).

GAL_DIMENSION_FLT_TO_INT (FLT) [Function-like macro]

Calculate the integer pixel position that the floating point FLT number belongs to. In the FITS format (and thus in Gnuastro), the center of each pixel is allocated on an integer (not it edge), so the pixel which hosts a floating point number cannot simply be found with internal type conversion.

void [Function]
gal_dimension_add_coords (*size_t* *c1, *size_t* *c2, *size_t* *out, *size_t* ndim)

For every dimension, add the coordinates in c1 with c2 and put the result into out. In other words, for dimension i run out[i]=c1[i]+c2[i];. Hence out may be equal to any one of c1 or c2.

size_t [Function]
gal_dimension_coord_to_index (*size_t* ndim, *size_t* *dsize, *size_t* *coord)

Return the index (counting from zero) from the coordinates in coord (counting from zero) assuming the dataset has ndim elements and the size of the dataset along each dimension is in the dsize array.

void [Function]
gal_dimension_index_to_coord (*size_t* index, *size_t* ndim, *size_t* *dsize, *size_t* *coord)

Fill in the coord array with the coordinates that correspond to index assuming the dataset has ndim elements and the size of the dataset along each dimension is in the dsize array. Note that both index and each value in coord are assumed to start from 0 (zero). Also that the space which coord points to must already be allocated before calling this function.

size_t [Function]
gal_dimension_dist_manhattan (*size_t* *a, *size_t* *b, *size_t* ndim)

Return the manhattan distance (see Wikipedia (https://en.wikipedia.org/wiki/Taxicab_geometry)) between the two coordinates a and b (each an array of ndim elements).

GAL_DIMENSION_NEIGHBOR_OP (index, ndim, dsize, connectivity, dinc, operation) [Function-like macro]

Parse the neighbors of the element located at index and do the requested operation on them. This is defined as a macro to allow easy definition of any operation on the neighbors of a given element without having to use loops within your source code (the loops are implemented by this macro). For an example of using this function, please see Section 10.4.2 [Library demo - inspecting neighbors], page 313. The input arguments to this function-like macro are described below:

index
: Distance of this element from the first element in the array on a contiguous patch of memory (starting from 0), see the discussion above.

ndim
: The number of dimensions associated with the contiugous patch of memory.

dsize
: The full array size along each dimension. This must be an array and is assumed to have the same number elements as ndim. See the discussion under the same element in Section 10.3.5.1 [Generic data container (gal_data_t)], page 245.

connectivity
: Most distant neighbors to consider. Depending on the number of dimensions, different neighbros may be defined for each element. This function-like macro distinguish between these different neighbors with this argument. It has a value between 1 (one) and ndim. For example in a 2D dataset, 4-connected neighbors have a connectivity of 1 and 8-connected neighbors have a connectivity of 2. Note that this is inclusive, so in this example, a connectivity of 2 will also include connectivity 1 neighbors.

dinc
: An array keeping the length necessary to increment along each dimension. You can make this array with the following function. Just don't forget to free the array after you are done with it:

```
size_t *dinc=gal_dimension_increment(ndim, dsize);
```

dinc depends on ndim and dsize, but it must be defined outside this function-like macro since it involves allocation to help in preformance.

operation
: Any C operation that you would like to do on the neighbor. This macro will provide you a nind variable that can be used as the index of the neighbor that is currently being studied. It is defined as 'size_t ndim;'. Note that operation will be repeated the number of times there is a neighbor for this element.

This macro works fully within its own {} block and except for the nind variable that shows the neighbor's index, all the variables within this macro's block start with gdn_.

10.3.7 Linked lists (list.h)

An array is a contiguous region of memory that is very efficient and easy to use for recording and later accessing any random element as fast as any other. This makes array the primary data container when you have many elements (for example an image which has millions of

pixels). One major problem with an array is that the number of elements that go into it must be known in advance and adding or removing an element will require a re-set of all the other elements. For example if you want to remove the 3rd element in a 1000 element array, all 997 subsequent elements have to pulled back by one position, the reverse will happen if you need to add an element.

In many contexts such situations never come up, for example you don't want to shift all the pixels in an image by one or two pixels from some random position in the image: their positions have scientific value. But in other contexts you will find your self frequently adding/removing an a-priori unknown of elements. Linked lists (or *lists* for short) are the data-container of choice in such situations. As in a chain, each *node* in a list is an independent C structure, keeping its own data along with pointer(s) to its immediate neighbor(s). Below, you can see one simple linked list node structure along with an ASCII art schematic of how we can use the `next` pointer to add any number of elements to the list that we want. By convension, a list is terminated when `next` is the `NULL` pointer.

```
struct list_float           /*     ---------    ---------           */
{                           /*     | Value |    | Value |           */
  float            value;   /*     |  ---  |    |  ---  |           */
  struct list_float *next;  /*     |  next-|--> |  next-|--> NULL   */
}                           /*     ---------    ---------           */
```

The schematic shows another great advantage of linked lists: it is very easy to add or remove/pop a node anywhere in the list. If you want to modify the first node, you just have to change one pointer. If it is in the middle, you just have to change two. You initially define a variable of this type with a `NULL` pointer as shown below:

```
struct list_float *mylist=NULL
```

then you use functions provided for that the respective type in the sections below to add elements to add or remove/pop an element from the list.

When you add an element to the list, it is conventionally added to the "top" of the list: the general list pointer will point to the newly created node, which will point to the previously created node and so on. So when you "pop" from the top of the list, you are actually retrieving the last value you put in and changing the list pointer to the next youngest node. This is thus known as a "last-in-first-out" list. This is the most efficient type of linked list (easier to implement and faster to process). Alternatively, you can add each newly created node at the end of the list. If you do that, you will get a "first-in-first-out" list. But that will force you to go through the whole list for each new element that is created (this will slow down the processing)[18].

The node example above creates the simplest kind of a list. We can define each node with two pointers to both the next and previous neighbors, this is called a "Doubly linked list". In general, lists are very powerful and simple constructs that can be very useful. But going into more detail would be out of the scope of this short introduction in this book. Wikipedia (`https://en.wikipedia.org/wiki/Linked_list`) has a nice and more thorough discussion of the various types of lists. To appreciate/use the beauty and elegance

[18] A better way to get a first-in-first-out is to first keep the data as last-in-first-out until they are all read. Afterwards, pop each node and immediately add it to the new list: practically reversing the last-in-first-out list to a first-in-first-out one. All the list types discussed in this chapter have a `_reverse` function.

of these powerful constructs even further, see Chapter 2 (Information Structures, in volume 1) of Donald Knuth's "The art of computer programming".

In this section we will review the functions and structures that are available in Gnuastro for working on lists. They differ by the type of data that each node can keep. For each linked-list node structure, we will first introduce the structure, then the functions for working on the structure. All these structures and functions are defined and declared in `gnuastro/list.h`.

10.3.7.1 List of strings

Probably one of the most common lists you will be using are lists of strings. They are the best tools when you are reading the user's inputs, or when adding comments to the output files. Below you can see Gnuastro's string list type and several functions to help in adding, removing/popping, reversing and freeing the list.

`gal_list_str_t` [Type (C `struct`)]

A single node in a list containing a string of characters.

```
typedef struct gal_list_str_t
{
  char *v;
  struct gal_list_str_t *next;
} gal_list_str_t;
```

`void` [Function]
`gal_list_str_add` (*gal_list_str_t* `**list`, *char* `*value`, *int* `allocate`)

Add a new node to the list of strings (`list`) and update it. The new node will contain the string `value`. If `allocate` is not zero, space will be allocated specifically for the string of the new node and the contents of `value` will be copied into it. This can be useful when your string may be changed later in the program, but you want your list to remain. Here is one short/simple example of initializing and adding elements to a string list:

```
gal_list_str_t *strlist;
gal_list_str_add(&strlist, "bottom of list.");
gal_list_str_add(&strlist, "second last element of list.");
```

`char *` [Function]
`gal_list_str_pop` (*gal_list_str_t* `**list`)

Pop the top element of `list`, change `list` to point to the next node in the list, and return the string that was in the popped node. If `*list==NULL`, then this function will also return a `NULL` pointer.

`size_t` [Function]
`gal_list_str_number` (*gal_list_str_t* `*list`)

Return the number of nodes in `list`.

`void` [Function]
`gal_list_str_print` (*gal_list_str_t* `*list`)

Print the strings within each node of `*list` on the standard output in the same order that they are stored. Each string is printed on one line. This function is mainly good for checking/debugging your program. For program outputs, its best to make

your own implementation with a better, more user-friendly, format. For example the following code snippet.

```
size_t i;
gal_list_str_t *tmp;
for(tmp=list; tmp!=NULL; tmp=tmp->next)
  printf("String %zu: %s\n", i, tmp->v);
```

void [Function]
gal_list_str_reverse (*gal_list_str_t* **list)

Reverse the order of the list such that the top node in the list before calling this function becomes the bottom node after it.

void [Function]
gal_list_str_free (*gal_list_str_t* *list, *int* freevalue)

Free every node in list. If freevalue is not zero, also free the string within the nodes.

10.3.7.2 List of int32_t

Signed integers are the best types when you are dealing with a positive or negative integers. The are generally useful in many contexts, for example when you want to keep the order of a series of states (each state stored as a given number in an enum for example). On many modern systems, int32_t is just an alias for int, so you can use them interchably. To make sure, check the size of int on your system:

gal_list_i32_t [Type (C struct)]

A single node in a list containing a 32-bit signed integer (see Section 4.4 [Numeric data types], page 64).

```
typedef struct gal_list_i32_t
{
  int32_t v;
  struct gal_list_i32_t *next;
} gal_list_i32_t;
```

void [Function]
gal_list_i32_add (*gal_list_i32_t* **list, *int32_t* value)

Add a new node (containing value) to the top of the list of int32_ts (uint32_t is equal to int on many modern systems), and update list. Here is one short example of initializing and adding elements to a string list:

```
gal_list_i32_t *i32list=NULL;
gal_list_i32_add(&i32list, 52);
gal_list_i32_add(&i32list, -4);
```

int32_t [Function]
gal_list_i32_pop (*gal_list_i32_t* **list)

Pop the top element of list and return the value. This function will also change list to point to the next node in the list. If *list==NULL, then this function will also return GAL_BLANK_INT32 (see Section 10.3.4 [Library blank values (blank.h)], page 243).

size_t [Function]
gal_list_i32_number (*gal_list_i32_t* *list)
Return the number of nodes in list.

void [Function]
gal_list_i32_print (*gal_list_i32_t* *list)
Print the integers within each node of *list on the standard output in the same order that they are stored. Each integer is printed on one line. This function is mainly good for checking/debugging your program. For program outputs, its best to make your own implementation with a better, more user-friendly format. For example the following code snippet. You can also modify it to print all values in one line, and etc, depending on the context of your program.

```
size_t i;
gal_list_i32_t *tmp;
for(tmp=list; tmp!=NULL; tmp=tmp->next)
  printf("String %zu: %s\n", i, tmp->v);
```

void [Function]
gal_list_i32_reverse (*gal_list_i32_t* **list)
Reverse the order of the list such that the top node in the list before calling this function becomes the bottom node after it.

int32_t * [Function]
gal_list_i32_to_array (*gal_list_i32_t* *list, *int* reverse, *size_t* *num)
Dynamically allocate an array and fill it with the values in list. The function will return a pointer to the allocated array and put the number of elements in the array into the num pointer. If reverse has a non-zero value, the array will be filled in the opposite order of elements in list. This function can be useful after you have finished reading an initially unknown number of values and want to put them in an array for easy random access.

void [Function]
gal_list_i32_free (*gal_list_i32_t* *list)
Free every node in list.

10.3.7.3 List of size_t

The size_t type is a unique type in C: as the name suggests it is defined to store sizes, or more accurately, the distances between memory locations. Hence it is always positive (an unsigned type) and it is directly related to the address-able spaces on the host system: on 32-bit and 64-bit systems it is an alias for uint32_t and uint64_t, respectively (see Section 4.4 [Numeric data types], page 64).

size_t is the default compiler type to index an array (recall that an array index in C is just a pointer increment of a given *size*). Since it is unsigned, its a great type for counting (where negative is not defined), you are always sure it will never exceed the system's (virtual) memory and since its name has the word "size" inside it, it provides a good level of documentation[19]. In Gnuastro, we do all counting and array indexing with

[19] So you know that a variable of this type is not used to store some generic state for example.

this type, so this list is very handy. As discussed above, `size_t` maps to different types on different machines, so a portable way to print them with `printf` is to use C99's `%zu` format.

`gal_list_sizet_t` [Type (C `struct`)]

A single node in a list containing a `size_t` value (which maps to `uint32_t` or `uint64_t` on 32-bit and 64-bit systems), see Section 4.4 [Numeric data types], page 64.

```
typedef struct gal_list_sizet_t
{
  size_t v;
  struct gal_list_sizet_t *next;
} gal_list_sizet_t;
```

`void` [Function]
`gal_list_sizet_add` (*gal_list_sizet_t* `**list`, *size_t* `value`)

Add a new node (containing `value`) to the top of the `list` of `size_t`s and update `list`. Here is one short example of initializing and adding elements to a string list:

```
gal_list_sizet_t *slist=NULL;
gal_list_sizet_add(&slist, 45493);
gal_list_sizet_add(&slist, 930484);
```

`sizet_t` [Function]
`gal_list_sizet_pop` (*gal_list_sizet_t* `**list`)

Pop the top element of `list` and return the value. This function will also change `list` to point to the next node in the list. If `*list==NULL`, then this function will also return `GAL_BLANK_SIZE_T` (see Section 10.3.4 [Library blank values (`blank.h`)], page 243).

`size_t` [Function]
`gal_list_sizet_number` (*gal_list_sizet_t* `*list`)

Return the number of nodes in `list`.

`void` [Function]
`gal_list_sizet_print` (*gal_list_sizet_t* `*list`)

Print the values within each node of `*list` on the standard output in the same order that they are stored. Each integer is printed on one line. This function is mainly good for checking/debugging your program. For program outputs, its best to make your own implementation with a better, more user-friendly format. For example, the following code snippet. You can also modify it to print all values in one line, and etc, depending on the context of your program.

```
size_t i;
gal_list_sizet_t *tmp;
for(tmp=list; tmp!=NULL; tmp=tmp->next)
  printf("String %zu: %zu\n", i, tmp->v);
```

`void` [Function]
`gal_list_sizet_reverse` (*gal_list_sizet_t* `**list`)

Reverse the order of the list such that the top node in the list before calling this function becomes the bottom node after it.

size_t * [Function]
gal_list_sizet_to_array (*gal_list_sizet_t* *list, *int* reverse, *size_t* *num)

Dynamically allocate an array and fill it with the values in list. The function will return a pointer to the allocated array and put the number of elements in the array into the num pointer. If reverse has a non-zero value, the array will be filled in the inverse of the order of elements in list. This function can be useful after you have finished reading an initially unknown number of values and want to put them in an array for easy random access.

void [Function]
gal_list_sizet_free (*gal_list_sizet_t* *list)

Free every node in list.

10.3.7.4 List of float

Single precision floating point numbers can accurately store real number until 7.2 decimals and only consume 4 bytes (32-bits) of memory, see Section 4.4 [Numeric data types], page 64. Since astronomical data rarely reach that level of precision, single precision floating points are the type of choice to keep and read data. However, when processing the data, it is best to use double precision floating points (since errors propagate).

gal_list_f32_t [Type (C struct)]

A single node in a list containing a 32-bit single precision float value: see Section 4.4 [Numeric data types], page 64.

```
typedef struct gal_list_f32_t
{
  float v;
  struct gal_list_f32_t *next;
} gal_list_f32_t;
```

void [Function]
gal_list_f32_add (*gal_list_f32_t* **list, *float* value)

Add a new node (containing value) to the top of the list of floats and update list. Here is one short example of initializing and adding elements to a string list:

```
gal_list_f32_t *flist=NULL;
gal_list_f32_add(&flist, 3.89);
gal_list_f32_add(&flist, 1.23e-20);
```

float [Function]
gal_list_f32_pop (*gal_list_f32_t* **list)

Pop the top element of list and return the value. This function will also change list to point to the next node in the list. If *list==NULL, then this function will return GAL_BLANK_FLOAT32 (NaN, see Section 10.3.4 [Library blank values (blank.h)], page 243).

size_t [Function]
gal_list_f32_number (*gal_list_f32_t* *list)

Return the number of nodes in list.

void [Function]
gal_list_f32_print (*gal_list_f32_t* *list)

Print the values within each node of *list on the standard output in the same order that they are stored. Each floating point number is printed on one line. This function is mainly good for checking/debugging your program. For program outputs, its best to make your own implementation with a better, more user-friendly format. For example, in the following code snippet. You can also modify it to print all values in one line, and etc, depending on the context of your program.

```
size_t i;
gal_list_f32_t *tmp;
for(tmp=list; tmp!=NULL; tmp=tmp->next)
  printf("Node %zu: %f\n", i, tmp->v);
```

void [Function]
gal_list_f32_reverse (*gal_list_f32_t* **list)

Reverse the order of the list such that the top node in the list before calling this function becomes the bottom node after it.

float * [Function]
gal_list_f32_to_array (*gal_list_f32_t* *list, *int* reverse, *size_t* *num)

Dynamically allocate an array and fill it with the values in list. The function will return a pointer to the allocated array and put the number of elements in the array into the num pointer. If reverse has a non-zero value, the array will be filled in the inverse of the order of elements in list. This function can be useful after you have finished reading an initially unknown number of values and want to put them in an array for easy random access.

void [Function]
gal_list_f32_free (*gal_list_f32_t* *list)

Free every node in list.

10.3.7.5 List of double

Double precision floating point numbers can accurately store real number until 15.9 decimals and consume 8 bytes (64-bits) of memory, see Section 4.4 [Numeric data types], page 64. This level of precision makes them very good for serious processing in the middle of a program's execution: in many cases, the propagation of errors will still be insignificant compared to actual obervational errors in a data set. But since they consume 8 bytes and more CPU processing power, they are often not the best choice for storing and transferring of data.

gal_list_f64_t [Type (C struct)]

A single node in a list containing a 64-bit double precision double value: see Section 4.4 [Numeric data types], page 64.

```
typedef struct gal_list_f64_t
{
  double v;
  struct gal_list_f64_t *next;
} gal_list_f64_t;
```

void [Function]
gal_list_f64_add (*gal_list_f64_t* **list, *double* value)

Add a new node (containing `value`) to the top of the `list` of `double`s and update `list`. Here is one short example of initializing and adding elements to a string list:

```
gal_list_f64_t *dlist=NULL;
gal_list_f64_add(&dlist, 3.8129395763193);
gal_list_f64_add(&dlist, 1.239378923931e-20);
```

double [Function]
gal_list_f64_pop (*gal_list_f64_t* **list)

Pop the top element of `list` and return the value. This function will also change `list` to point to the next node in the list. If `*list==NULL`, then this function will return `GAL_BLANK_FLOAT64` (NaN, see Section 10.3.4 [Library blank values (`blank.h`)], page 243).

size_t [Function]
gal_list_f64_number (*gal_list_f64_t* *list)

Return the number of nodes in `list`.

void [Function]
gal_list_f64_print (*gal_list_f64_t* *list)

Print the values within each node of `*list` on the standard output in the same order that they are stored. Each floating point number is printed on one line. This function is mainly good for checking/debugging your program. For program outputs, its best to make your own implementation with a better, more user-friendly format. For example, in the following code snippet. You can also modify it to print all values in one line, and etc, depending on the context of your program.

```
size_t i;
gal_list_f64_t *tmp;
for(tmp=list; tmp!=NULL; tmp=tmp->next)
  printf("Node %zu: %f\n", i, tmp->v);
```

void [Function]
gal_list_f64_reverse (*gal_list_f64_t* **list)

Reverse the order of the list such that the top node in the list before calling this function becomes the bottom node after it.

double * [Function]
gal_list_f64_to_array (*gal_list_f64_t* *list, *int* reverse, *size_t* *num)

Dynamically allocate an array and fill it with the values in `list`. The function will return a pointer to the allocated array and put the number of elements in the array into the `num` pointer. If `reverse` has a non-zero value, the array will be filled in the inverse of the order of elements in `list`. This function can be useful after you have finished reading an initially unknown number of values and want to put them in an array for easy random access.

void [Function]
gal_list_f64_free (*gal_list_f64_t* *list)

Free every node in `list`.

10.3.7.6 List of `void *`

In C, `void *` is the most generic pointer. Usually pointers are associated with the type of content they point to. For example `int *` means a pointer to an integer. This ancillary information about the contents of the memory location is very useful for the compiler, catching bad errors and also documentation (it helps the reader see what the address in memory actually contains). However, `void *` is just a raw address (pointer), it contains no information on the contents it points to.

These properties make the `void *` very useful when you want to treat the contents of an address in differnt ways. You can use the `void *` list defined in this section and its function on any kind of data: for example you can use it to keep a list of custom data structures that you have built for your own separate program. Each node in the list can keep anything and this gives you great versatily. But in using `void *`, please beware that "with great power comes great responsibilty".

`gal_list_void_t` [Type (C `struct`)]

A single node in a list containing a `void *` pointer.

```
typedef struct gal_list_void_t
{
  void *v;
  struct gal_list_void_t *next;
} gal_list_void_t;
```

`void` [Function]
`gal_list_void_add` (*gal_list_void_t* `**list`, *void* `*value`)

Add a new node (containing `value`) to the top of the `list` of `void *`s and update `list`. Here is one short example of initializing and adding elements to a string list:

```
gal_list_void_t *vlist=NULL;
gal_list_f64_add(&vlist, some_pointer);
gal_list_f64_add(&vlist, another_pointer);
```

`void *` [Function]
`gal_list_void_pop` (*gal_list_void_t* `**list`)

Pop the top element of `list` and return the value. This function will also change `list` to point to the next node in the list. If `*list==NULL`, then this function will return `NULL`.

`size_t` [Function]
`gal_list_void_number` (*gal_list_void_t* `*list`)

Return the number of nodes in `list`.

`void` [Function]
`gal_list_void_reverse` (*gal_list_void_t* `**list`)

Reverse the order of the list such that the top node in the list before calling this function becomes the bottom node after it.

`void` [Function]
`gal_list_void_free` (*gal_list_void_t* `*list`)

Free every node in `list`.

10.3.7.7 Ordered list of size_t

Positions/sizes in a dataset are conventionally in the size_t type (see Section 10.3.7.3 [List of size_t], page 259) and it sometimes occurs that you want to parse and read the values in a specific order. For example you want to start from one pixel and add pixels to the list based on their distance to that pixel. So that ever time you pop an element from the list, you know it is the nearest that has not yet been studied. The gal_list_osizet_t type and its functions in this section are designed to facilitate such operations.

gal_list_osizet_t [Type (C struct)]

Each node in this singly-linked list contains a size_t value and a floating point value. The floating point value is used as a reference to add new nodes in a sorted manner. At any moment, the first popped node in this list will have the smallest tosort value, and subsequent nodes will have larger to values.

```
typedef struct gal_list_osizet_t
{
  size_t v;                       /* The actual value. */
  float  s;                       /* The parameter to sort by. */
  struct gal_list_osizet_t *next;
} gal_list_osizet_t;
```

void [Function]
gal_list_osizet_add (*gal_list_osizet_t* **list, *size_t* value, *float* tosort)

Allocate space for a new node in list, and store value and tosort into it. The new node will not necessarily be at the "top" of the list. If *list!=NULL, then the tosort values of existing nodes is inspected and the given node is placed in the list such that the top element (which is popped with gal_list_osizet_pop) has the smallest tosort value.

size_t [Function]
gal_list_osizet_pop (*gal_list_osizet_t* **list, *float* *sortvalue)

Pop a node from the top of list, return the node's value and put its sort value in the space that sortvalue points to. This function will also free the allocated space for the popped node and after this function, list will point to the next node (which has a larger tosort element).

void [Function]
gal_list_osizet_to_sizet_free (*gal_list_osizet_t* *in, *gal_list_sizet_t* **out)

Convert the orderd list of size_ts into an ordinary size_t linked list. This can be useful when all the elements have been added and you just need to pop-out elements and don't care about the sorting values any more. After the conversion is done, this function will free the input list. Note that the out list doesn't have to be empty. If it already contains some nodes, the new nodes will be added ontop of them.

10.3.7.8 Doubly linked ordered list of size_t

An ordered list of indexs is required in many contexts, one example was discussed at the beginning of Section 10.3.7.7 [Ordered list of size_t], page 265. But the list that was introduced there only has one point of entry: you can always only parse the list from

smallest to largest. In this section, the doubly-linked `gal_list_dosizet_t` node is defined which will allow us to parse the values in ascending or descending order.

`gal_list_dosizet_t` [Type (C `struct`)]

Doubly-linked, ordered `size_t` list node structure. Each node in this Doubly-linked list contains a `size_t` value and a floating point value. The floating point value is used as a reference to add new nodes in a sorted manner. In the functions here, this linked list can be pointed to by two pointers (largest and smallest) with the following format:

```
            largest pointer
            |
   NULL <-- (v0,s0) <--> (v1,s1) <--> ... (vn,sn) --> NULL
                                          |
                           smallest pointer
```

At any moment, the two pointers will point to the nodes containing the "largest" and "smallest" values and the rest of the nodes will be sorted. This is useful when an unknown number of nodes are being added continuously and during the operations it is important to have the nodes in a sorted format.

```
typedef struct gal_list_dosizet_t
{
  size_t v;                          /* The actual value. */
  float s;                           /* The parameter to sort by. */
  struct gal_list_dosizet_t *prev;
  struct gal_list_dosizet_t *next;
} gal_list_dosizet_t;
```

`void` [Function]
`gal_list_dosizet_add` (*gal_list_dosizet_t* `**largest`, *gal_list_dosizet_t* `**smallest`, *size_t* `value`, *float* `tosort`)

Allocate space for a new node in `list`, and store `value` and `tosort` into it. If the list is empty, both `largest` and `smallest` must be `NULL`.

`size_t` [Function]
`gal_list_dosizet_pop_smallest` (*gal_list_dosizet_t* `**largest`, *gal_list_dosizet_t* `**smallest`, *float* `tosort`)

Pop the value with the smallest reference from the doubly linked list and store the reference into the space pointed to by `tosort`. Note that even though only the smallest pointer will be popped, when there was only one node in the list, the `largest` pointer also has to change, so we need both.

`void` [Function]
`gal_list_dosizet_print` (*gal_list_dosizet_t* `*largest`, *gal_list_dosizet_t* `*smallest`)

Print the largest and smallest values sequentially until the list is parsed.

`void` [Function]
`gal_list_dosizet_to_sizet` (*gal_list_dosizet_t* `*in`, *gal_list_sizet_t* `**out`)

Convert the doubly linked, ordered `size_t` list into a singly-linked list of `size_t`.

void [Function]
`gal_list_dosizet_free` (*gal_list_dosizet_t* *largest)

Free the doubly linked, ordered `sizet_t` list.

10.3.7.9 List of `gal_data_t`

Gnuastro's generic data container has a `next` element which enables it to be used as a singly-linked list (see Section 10.3.5.1 [Generic data container (`gal_data_t`)], page 245). The ability to connect the different data containers offers great advantages. For example each column in a table in an independent dataset: with its own name, units, numeric data type (see Section 4.4 [Numeric data types], page 64). Another application is in Tessellating an input dataset into separate tiles or only studying particular regions, or tiles, of a larger dataset (see Section 4.6 [Tessellation], page 72, and Section 10.3.13 [Tessellation library (`tile.h`)], page 288). Each independent tile over the dataset can be connected to the others as a linked list and thus any number of tiles can be represented with one variable.

void [Function]
`gal_list_data_add` (*gal_data_t* **list, *gal_data_t* *newnode)

Add an already allocated dataset (`newnode`) to top of `list`. In this example multiple images are linked together as a list:

```
size_t minmapsize=-1;
gal_data_t *tmp, *list=NULL;
tmp = gal_fits_img_read("file1.fits", "1", minmapsize);
gal_list_data_add( &list, tmp );
tmp = gal_fits_img_read("file2.fits", "1", minmapsize);
gal_list_data_add( &list, tmp );
```

void [Function]
`gal_list_data_add_alloc` (*gal_data_t* **list, *void* *array, *uint8_t* type, *size_t* ndim, *size_t* *dsize, *struct wcsprm* *wcs, *int* clear, *size_t* minmapsize, *char* *name, *char* *unit, *char* *comment)

Allocate a new dataset (with `gal_data_alloc` in Section 10.3.5.2 [Dataset size and allocation], page 250) and put it as the first element of `list`. Note that if this is the first node to be added to the list, `list` must be `NULL`.

gal_data_t * [Function]
`gal_list_data_pop` (*gal_data_t* **list)

Pop the top node from `list` and return it.

void [Function]
`gal_list_data_reverse` (*gal_data_t* **list)

Reverse the order of the list such that the top node in the list before calling this function becomes the bottom node after it.

size_t [Function]
`gal_list_data_number` (*gal_data_t* *list)

Return the number of nodes in `list`.

void [Function]
`gal_list_data_free` (*gal_data_t* *list)

Free all the datasets in `list` along with all the allocated spaces in each.

10.3.8 FITS files (`fits.h`)

The FITS format is the most common format to store data (images and tables) in astronomy. The CFITSIO library already provides a very good low-level collection of functions for manipulating FITS data. The low-level nature of CFITSIO is defined for versatility and portability. As a result, even a simple a basic operation, like reading an image or table column into memory, will require a special sequence of CFITSIO function calls which can be inconvenient and buggy to manage in separate locations. Therefore Gnuastro library provides wrappers for CFITSIO functions to make it much easier to read/write/modify FITS file data, header keywords and extensions. Hence, if you feel these functions don't exactly do what you want, we strongly recommend reading the CFITSIO manual to use its great features directly.

All the functions and macros introduced in this section are declared in `gnuastro/fits.h`. When you include this header, you are also including CFITSIO's `fitsio.h` header. So you don't need to explicitly include `fitsio.h` anymore and can freely use any of its macros or functions in your code along with those discussed here.

10.3.8.1 FITS Macros, errors and filenames

Some general constructs provided by Gnuastro's FITS handling functions are discussed here. In particular there are several useful functions about FITS file names.

`GAL_FITS_MAX_NDIM` [Macro]

The maximum number of dimensions a dataset can have in FITS format, according to the FITS standard this is 999.

`void` [Function]
`gal_fits_io_error` (*int* `status`, *char* `*message`)

If `status` is non-zero, this function will print the CFITSIO error message corresponding to status, print `message` (optional) in the next line and abort the program. If `message==NULL`, it will print a default string after the CFITSIO error.

`int` [Function]
`gal_fits_name_is_fits` (*char* `*name`)

If the `name` is an acceptable CFITSIO FITS filename return `1` (one), otherwise return `0` (zero). The currently acceptable FITS suffixes are `.fits`, `.fits.gz`, `.fits.Z`, `.imh`, `.fits.fz`. IMH is the IRAF format which is acceptable to CFITSIO.

`int` [Function]
`gal_fits_suffix_is_fits` (*char* `*suffix`)

Similar to `gal_fits_name_is_fits`, but only for the suffix. The suffix doesn't have to start with '`.`': this function will return `1` (one) for both `fits` and `.fits`.

`char *` [Function]
`gal_fits_name_save_as_string` (*char* `*filename`, *char* `*hdu`)

If the name is a FITS name, then put a `(hdu: ...)` after it and return the string. If it isn't a FITS file, just print the name. Note that the space is allocated. This function is useful when you want to report a random file to the user which may be FITS or not (for a FITS file, simply the filename is not enough, the HDU is also necessary).

10.3.8.2 CFITSIO and Gnuastro types

Both Gnuastro and CFITSIO have special identifiers for each type that they accept. Gnuastro's type identifiers are fully described in Section 10.3.3 [Library data types (`type.h`)], page 238, and are usable for all kinds of datasets (images, table columns and etc) as part of Gnuastro's Section 10.3.5.1 [Generic data container (`gal_data_t`)], page 245. However, following the FITS standard, CFITSIO has different identifiers for images and tables. Following CFITSIO's own convention, we will use `bitpix` for image type identifiers and `datatype` for its internal identifiers (and mainly used in tables). The functions introduced in this section can be used to convert between CFITSIO and Gnuastro's type identifiers.

One important issue to consider is that CFITSIO's types are not fixed width (for example `long` may be 32-bits or 64-bits on different systems). However, Gnuastro's types are defined by their width. These functions will use information on the host system to do the proper conversion, so it strongly recommended to use these functions for portability of your code and not to assume a fixed correspondance between CFITSIO and Gnuastro's types.

`uint8_t` [Function]
`gal_fits_bitpix_to_type` (*int* `bitpix`)
Return the Gnuastro type identifier that corresponds to CFITSIO's `bitpix` on this system.

`int` [Function]
`gal_fits_type_to_bitpix` (*uint8_t* `type`)
Return the CFITSIO `bitpix` value that corresponds to Gnuastro's `type`.

`char` [Function]
`gal_fits_type_to_bin_tform` (*uint8_t* `type`)
Return the FITS standard binary table `TFORM` character that corresponds to Gnuastro's `type`.

`int` [Function]
`gal_fits_type_to_datatype` (*uint8_t* `type`)
Return the CFITSIO `datatype` that corresponds to Gnuastro's `type` on this machine.

`uint8_t` [Function]
`gal_fits_datatype_to_type` (*int* `datatype`, *int* `is_table_column`)
Return Gnuastro's type identifier that corresponds to the CFITSIO `datatype`. Note that when dealing with CFITSIO's `TLONG`, the fixed width type differs between tables and images. So if the corresponding dataset is a table column, put a non-zero value into `is_table_column`.

10.3.8.3 FITS HDUs

A FITS file can contain multiple HDUs/extensions. The functions in this section can be used to get basic information about the extensions or open them. Note that `fitsfile` is defined in CFITSIO's `fitsio.h` which is automatically included by Gnuastro's `gnuastro/fits.h`.

fitsfile * [Function]
gal_fits_open_to_write (*char* *filename)

If filename exists, open it and return the fitsfile pointer that corresponds to it. If filename doesn't exist, the file will be created which contains a blank first extension and the pointer to its next extension will be returned.

size_t [Function]
gal_fits_hdu_num (*char* *filename)

Return the number of HDUs/extensions in filename.

int [Function]
gal_fits_hdu_format (*char* *filename, *char* *hdu)

Return the format of the HDU as one of CFITSIO's recognized macros: IMAGE_HDU, ASCII_TBL, or BINARY_TBL.

fitsfile * [Function]
gal_fits_hdu_open (*char* *filename, *char* *hdu, *int* iomode)

Open the HDU/extension hdu from filename and return a pointer to CFITSIO's fitsfile. iomode determines how the FITS file will be opened using CFITSIO's macros: READONLY or READWRITE.

The string in hdu will be appended to filename in square brackets so CFITSIO only opens this extension. You can use any formatting for the hdu that is acceptable to CFITSIO. See the description under --hdu in Section 4.1.2.1 [Input/Output options], page 52, for more.

fitsfile * [Function]
gal_fits_hdu_open_format (*char* *filename, *char* *hdu, *int* img0_tab1)

Open (in read-only format) the hdu HDU/extension of filename as an image or table. When img0_tab1 is 0(zero) but the HDU is a table, this function will abort with an error. It will also abort with an error when img0_tab1 is 1 (one), but the HDU is an image.

A FITS HDU may contain both tables or images. When your program needs one of these formats, you can call this function so if the user provided the wrong HDU/file, it will abort and inform the user that the file/HDU is has the wrong format.

10.3.8.4 FITS header keywords

Each FITS extension/HDU contains a raw dataset which can either be a table or an image along with some header keywords. The keywords can be used to store meta-data about the actual dataset. The functions in this section describe Gnuastro's high-level functions for reading and writing FITS keywords. Similar to all Gnuastro's FITS-related functions, these functions are all wrappers for CFITSIO's low-level functions.

The necessary meta-data (header keywords) for a particular dataset are commonly numerous, it is much more efficient to list them in one variable and call the reading/writing functions once. Hence the functions in this section use linked lists, a thorough introduction to them is given in Section 10.3.7 [Linked lists (list.h)], page 255. To reading FITS keywords, these functions use a list of Gnuastro's generic dataset format that is discussed in Section 10.3.7.9 [List of gal_data_t], page 267. To write FITS keywords we define the gal_fits_list_key_t node that is defined below.

`gal_fits_list_key_t` [Type (C `struct`)]

Structure for writing FITS keywords. This structure is used for one keyword and you don't need to set all elements. With the `next` element, you can link it to another keyword thus creating a linked list to add any number of keywords easily and at any step during your program (see Section 10.3.7 [Linked lists (`list.h`)], page 255, for an introduction on lists). See the functions below for adding elements to the list.

```
typedef struct gal_fits_list_key_t
{
  int                        kfree;   /* ==1, free name.      */
  int                        vfree;   /* ==1, free value.     */
  int                        cfree;   /* ==1, free comment.   */
  uint8_t                     type;   /* Keyword value type.  */
  char                    *keyname;   /* Keyword Name.        */
  void                      *value;   /* Keyword value.       */
  char                    *comment;   /* Keyword comment.     */
  char                       *unit;   /* Keyword unit.        */
  struct gal_fits_list_key_t *next;   /* Next keyword.        */
} gal_fits_list_key_t;
```

`void` [Function]
`gal_fits_key_clean_str_value` (*char* `*string`)

Remove the sigle quotes and possible extra spaces around the keyword values that CFITSIO returns when reading a string keyword. CFITSIO doesn't remove the two single quotes around the string value of a keyword. Hence the strings it reads are like: `'value '`, or `'some_very_long_value'`. To use the value during your processing, it is commonly necessary to remove the single quotes (and possible extra spaces). This function will do this within the allocated space of the string.

`void` [Function]
`gal_fits_key_read_from_ptr` (*fitsfile* `*fptr`, *gal_data_t* `*keysll`, *int* `readcomment`, *int* `readunit`)

Read the list of keyword values from a FITS pointer. The input should be a linked list of Gnuastro's generic data container (`gal_data_t`). Before calling this function, you just have to set the `name` and desired `type` values of each node in the list to the keyword you want it to keep the value of. The given `name` value will be directly passed to CFITSIO to read the desired keyword name. This function will allocate space to keep the value. If `readcomment` and `readunit` are non-zero, this function will also try to read the possible comments and units of the keyword.

Here is one example of using this function:

```
/* Allocate an array of datasets. */
gal_data_t *keysll=gal_data_array_calloc(N);

/* Use the array as a list.*/
for(i=0;i<N-2;++i) keysll[i].next=keysll[i+1];

/* Fill the datasets with a `name' and a `type'. */
keysll[0].name="NAME1";     keysll[0].type=GAL_TYPE_INT32;
```

```
keysll[1].name="NAME2";     keysll[1].type=GAL_TYPE_STRING;
...
...

/* Call this function. */
gal_fits_key_read_from_ptr(fptr, keysll, 0, 0);

/* Use the values as you like... */

/* Free all the allocated spaces. */
gal_data_array_free(keysll, N, 1);
```

If the `array` pointer of each keyword's dataset is not `NULL`, then it is assumed that the space to keep the value has already been allocated. If it is `NULL`, space will be allocated for the value by this function.

Strings need special consideration: the reason is that generally, `gal_data_t` needs to also allow for array of strings (as it supports arrays of integers for example). Hence when reading a string value, two allocations may be done by this function (one if `array!=NULL`).

Therefore, when using the values of strings after this function, `keysll[i].array` must be interpretted as `char **`: one allocation for the pointer, one for the actual characters. If you use something like the example, above you don't have to worry about the freeing, `gal_data_array_free` will free both allocations. So to read a string, one easy way would be the following:

```
char *str, **strarray;
strarr = keysll[i].array;
str    = strarray[0];
```

If CFITSIO is unable to read a keyword for any reason the `status` element of the respective `gal_data_t` will be non-zero. If it is zero, then the keyword was found and succesfully read. Otherwise, it is a CFITSIO status value. You can use CFITSIO's error reporting tools or `gal_fits_io_error` (see Section 10.3.8.1 [FITS Macros, errors and filenames], page 268) for reporting the reason of the failure. A tip: when the keyword doesn't exist, CFITSIO's status value will be `KEY_NO_EXIST`.

CFITSIO will start searching for the keywords from the last place in the header that it searched for a keyword. So it is much more efficient if the order that you ask for keywords is based on the order they are stored in the header.

`void` [Function]
`gal_fits_key_read` (*char* `*filename`, *char* `*hdu`, *gal_data_t* `*keysll`, *int* `readcomment`, *int* `readunit`)

Same as `gal_fits_read_keywords_fptr` (see above), but accepts the filename and HDU as input instead of an already opened CFITSIO `fitsfile` pointer.

void [Function]
`gal_fits_key_list_add` (*gal_fits_list_key_t* `**list`, *uint8_t* `type`, *char* `*keyname`, *int* `kfree`, *void* `*value`, *int* `vfree`, *char* `*comment`, *int* `cfree`, *char* `*unit`)

Add on keyword to the top of list of header keywords that need to be written into a FITS file. In the end, the keywords will have to be freed, so it is important to know before hand if they were allocated or not (hence the presence of the arguments ending in `free`). If the space for the respective element is not allocated, set these arguments to `0` (zero).

Important note for strings: the value should be the pointer to the string its-self (`char *`), not a pointer to a pointer (`char **`).

void [Function]
`gal_fits_key_list_add_end` (*gal_fits_list_key_t* `**list`, *uint8_t* `type`, *char* `*keyname`, *int* `kfree`, *void* `*value`, *int* `vfree`, *char* `*comment`, *int* `cfree`, *char* `*unit`)

Similar to `gal_fits_key_list_add` (see above) but add the keyword to the end of the list. Use this function if you want the keywords to be written in the same order that you add nodes to the list of keywords.

void [Function]
`gal_fits_key_write_filename` (*char* `*keynamebase`, *char* `*filename`, *gal_fits_list_key_t* `**list`)

Put `filename` into the `gal_fits_list_key_t` list (possibly broken up into multiple keywords) to later write into a HDU header. The `keynamebase` string will be appended with a `_N` (N>0) and used as the keyword name.

The FITS standard sets a maximum length for the value of a keyword. This creates problems with file names (which include directories). Because file names/addresses can become very long. Therefore, when `filename` is longer than the maximum length of a FITS keyword value, this function will break it into several keywords (breaking up the string on directory separators).

void [Function]
`gal_fits_key_write_wcsstr` (*fitsfile* `*fptr`, *char* `*wcsstr`, *int* `nkeyrec`)

Write the WCS header string (produced with WCSLIB's `wcshdo` function) into the CFITSIO `fitsfile` pointer. `nkeyrec` is the number of FITS header keywords in `wcsstr`. This function will put a few blank keyword lines along with a comment `WCS information` before writing each keyword record.

void [Function]
`gal_fits_key_write` (*fitsfile* `*fptr`, *gal_fits_list_key_t* `**keylist`)

Write the list of keywords in `keylist` into the CFITSIO `fitsfile`.

void [Function]
`gal_fits_key_write_version` (*fitsfile* `*fptr`, *gal_fits_list_key_t* `*headers`, *char* `*program_name`)

Write or update (all the) keyword(s) in `headers` into the FITS pointer, but also the date, name of your program (`program_name`), along with the versions of CFITSIO,

WCSLIB (when available), GSL, Gnuastro, and (the possible) commit information into the header as described in Section 4.9 [Output headers], page 78.

Since the data processing depends on the versions of the libraries you have used, it is strongly recommended to include this information in every FITS output. `gal_fits_img_write` and `gal_fits_tab_write` will automatically use this function.

10.3.8.5 FITS arrays (images)

Images (or multi-dimensional arrays in general) are one of the common data formats that is stored in FITS files. Only one image may be stored in each FITS HDU/extension. The functions described here can be used to get the information of, read, or write images in FITS files.

`void` [Function]
`gal_fits_img_info` (*fitsfile* `*fptr`, *int* `*type`, *size_t* `*ndim`, *size_t* `**dsize`)

Read the type (see Section 10.3.3 [Library data types (`type.h`)], page 238), number of dimensions, and size of the array along each dimension of the CFITSIO `fitsfile` into the `type`, `ndim`, and `dsize` pointers respectively.

`gal_data_t *` [Function]
`gal_fits_img_read` (*char* `*filename`, *char* `*hdu`, *size_t* `minmapsize`)

Read the contents of the `hdu` extension/HDU of `filename` into a Gnuastro generic data container (see Section 10.3.5.1 [Generic data container (`gal_data_t`)], page 245) and return it. If the necessary space is larger than `minmapsize`, then don't keep the data in RAM, but in a file on the HDD/SSD. For more on `minmapsize` see the description under the same name in Section 10.3.5.1 [Generic data container (`gal_data_t`)], page 245.

`gal_data_t *` [Function]
`gal_fits_img_read_to_type` (*char* `*inputname`, *char* `*inhdu`, *uint8_t* `type`, *size_t* `minmapsize`)

Read the contents of the `hdu` extension/HDU of `filename` into a Gnuastro generic data container (see Section 10.3.5.1 [Generic data container (`gal_data_t`)], page 245) of type `type` and return it.

This is just a wrapper around `gal_fits_img_read` (to read the image/array of any type) and `gal_data_copy_to_new_type_free` (to convert it to `type` and free the initially read dataset).

`gal_data_t *` [Function]
`gal_fits_img_read_kernel` (*char* `*filename`, *char* `*hdu`, *size_t* `minmapsize`)

Read the `hdu` of `filename` as a convolution kernel. A convolution kernel must have an odd size along all dimensions, it must not have blank (NaN in floating point types) values and must be flipped around the center to make the proper convolution (see Section 6.3.1.1 [Convolution process], page 118). If there are blank values, this function will change the blank values to `0.0`. If the input image doesn't have the other two requirements, this function will abort with an error describing the condition to the user. The finally returned dataset will have a `float32` type.

fitsfile * [Function]
`gal_fits_img_write_to_ptr` (*gal_data_t* `*input`, *char* `*filename`)

Write the `input` dataset into a FITS file named `filename` and return the corresponding CFITSIO `fitsfile` pointer. This function won't close `fitsfile`, so you can still add other extensions to it after this function or make other modifications.

void [Function]
`gal_fits_img_write` (*gal_data_t* `*data`, *char* `*filename`, *gal_fits_list_key_t* `*headers`, *char* `*program_string`)

Write the `input` dataset into the FITS file named `filename`. Also add the `headers` keywords to the newly created HDU/extension it along with your program's name (`program_string`).

void [Function]
`gal_fits_img_write_to_type` (*gal_data_t* `*data`, *char* `*filename`, *gal_fits_list_key_t* `*headers`, *char* `*program_string`, *int* `type`)

Convert the `input` dataset into `type`, then write it into the FITS file named `filename`. Also add the `headers` keywords to the newly created HDU/extension along with your program's name (`program_string`). After the FITS file is written, this function will free the copied dataset (with type `type`) from memory.

This is just a wrapper for the `gal_data_copy_to_new_type` and `gal_fits_img_write` functions.

void [Function]
`gal_fits_img_write_corr_wcs_str` (*gal_data_t* `*data`, *char* `*filename`, *char* `*wcsstr`, *int* `nkeyrec`, *double* `*crpix`, *gal_fits_list_key_t* `*headers`, *char* `*program_string`)

Write the `input` dataset into `filename` using the `wcsstr` while correcting the `CRPIX` values.

This function is mainly useful when you want to make FITS files in parallel (from one main WCS structure, with just differing CRPIX). This can happen in the following cases for example:

- When a large number of FITS images (with WCS) need to be created in parallel, it can be much more efficient to write the header's WCS keywords once at first, write them in the FITS file, then just correct the CRPIX values.
- WCSLIB's header writing function is not thread safe. So when writing FITS images in parallel, we can't write the header keywords in each thread.

10.3.8.6 FITS tables

Tables are one of the common formats of data that is stored in FITS files. Only one table may be stored in each FITS HDU/extension, but each table column must be viewed as a different dataset (with its own name, units and numeric data type for example). The only constraint of the column datasets in a table is that they must be one-dimensional and have the same number of elements as the other columns. The functions described here can be used to get the information of, read, or write columns into FITS tables.

void [Function]
gal_fits_tab_size (*fitsfile* *fitsptr, *size_t* *nrows, *size_t* *ncols)
Read the number of rows and columns in the table within CFITSIO's fitsptr.

int [Function]
gal_fits_tab_format (*fitsfile* *fitsptr)
Return the format of the FITS table contained in CFITSIO's fitsptr. Recall that FITS tables can be in binary or ASCII formats. This function will return GAL_TABLE_FORMAT_AFITS or GAL_TABLE_FORMAT_BFITS (defined in Section 10.3.11 [Table input output (table.h)], page 281). If the fitsptr is not a table, this function will abort the program with an error message informing the user of the problem.

gal_data_t * [Function]
gal_fits_tab_info (*char* *filename, *char* *hdu, *size_t* *numcols, *size_t* *numrows, *int* *tableformat)
Store the information of each column in hdu of filename into an array of data structures with numcols elements (one data structure for each column) see Section 10.3.5.3 [Arrays of datasets], page 252. The total number of rows in the table is also put into the memory that numrows points to. The format of the table (e.g., FITS binary or ASCII table) will be put in tableformat (macros defined in Section 10.3.11 [Table input output (table.h)], page 281).

Note that other than the character strings (column name, units and comments), nothing in the data structure(s) will be allocated by this function for the actual data (e.g., the array or dsize elements). This function is just for column information. This is a low-level function particular to reading tables in FITS format. It is recommended to use gal_table_info to be generic (see Section 10.3.11 [Table input output (table.h)], page 281).

gal_data_t * [Function]
gal_fits_tab_read (*char* *filename, *char* *hdu, *size_t* numrows, *gal_data_t* *colinfo, *gal_list_sizet_t* *indexll, *int* minmapsize)
Read the columns given in the list indexll from a FITS table into a linked list of data structures, see Section 10.3.7.3 [List of size_t], page 259, and Section 10.3.7.9 [List of gal_data_t], page 267. If the necessary space for each column is larger than minmapsize, don't keep it in the RAM, but in a file in the HDD/SSD, see the description under the same name in Section 10.3.5.1 [Generic data container (gal_data_t)], page 245.

Note that this is a low-level function, so the output data linked list is the inverse of the input indexs linked list. It is recommended to use gal_table_read for generic reading of tables, see Section 10.3.11 [Table input output (table.h)], page 281.

void [Function]
gal_fits_tab_write (*gal_data_t* *cols, *gal_list_str_t* *comments, *int* tableformat, *char* *filename, *int* dontdelete)
Write the list of datasets in cols (see Section 10.3.7.9 [List of gal_data_t], page 267) as separate columns in a FITS table in filename. If filename already exists and dontdelete is non-zero, then this function will abort with an error and will not write over the existing file. The format of the table (ASCII or binary) may be specified with

the `tableformat` (see Section 10.3.11 [Table input output (`table.h`)], page 281). If `comments!=NULL`, each node of the list of strings will be written as a `COMMENT` keywords in the output FITS file (see Section 10.3.7.1 [List of strings], page 257.

10.3.9 World Coordinate System (`wcs.h`)

The FITS standard defines the world coordinate system (WCS) as a mechanism to associate physical values to positions within a dataset. For example, it can be used to convert pixel coordinates in an image to celestial coordinates like the right ascension and declination. The functions in this section are mainly just wrappers over CFITSIO, WCSLIB and GSL library functions to help in common applications.

`struct wcsprm *` [Function]
`gal_wcs_read_fitsptr` (*fitsfile* `*fptr`, *size_t* `hstartwcs`, *size_t* `hendwcs`, *int* `*nwcs`)

[**Not thread-safe**] Return the WCSLIB `wcsprm` structure that is read from the CFITSIO `fptr` pointer to an opened FITS file. Also put the number of coordinate representations found into the space that `nwcs` points to. To read the WCS structure directly from a filename, see `gal_wcs_read` below. After processing has finished, you can free the returned structure with WCSLIB's `wcsvfree` keyword:

```
status = wcsvfree(&nwcs,&wcs);
```

If you don't want to search the full FITS header for WCS-related FITS keywords (for example due to conflicting keywords), but only a specific range of the header keywords you can use the `hstartwcs` and `hendwcs` arguments to specify the keyword number range (counting from zero). If `hendwcs` is larger than `hstartwcs`, then only keywords in the given range will be checked. Hence, to ignore this feature (and search the full FITS header), give both these arguments the same value.

If the WCS information couldn't be read from the FITS file, this function will return a `NULL` pointer and put a zero in `nwcs`. A WCSLIB error message will also be printed in `stderr` if there was an error.

This function is just a wrapper over WCSLIB's `wcspih` function which is not thread-safe. Therefore, be sure to not call this function simultaneously (over multiple threads).

`struct wcsprm *` [Function]
`gal_wcs_read` (*char* `*filename`, *char* `*hdu`, *size_t* `hstartwcs`, *size_t* `hendwcs`, *int* `*nwcs`)

[**Not thread-safe**] Return the WCSLIB structure that is read from the HDU/extension `hdu` of the file `filename`. Also put the number of coordinate representations found into the space that `nwcs` points to. Please see `gal_wcs_read_fitsptr` for more.

`struct wcsprm *` [Function]
`gal_wcs_copy` (*struct wcsprm* `*wcs`)

Return a fully allocated (independent) copy of `wcs`.

`void` [Function]
`gal_wcs_on_tile` (*gal_data_t* `*tile`)

Create a WCSLIB `wcsprm` structure for `tile` using WCS parameters of the tile's allocated block dataset, see Section 10.3.13 [Tessellation library (`tile.h`)], page 288,

for the definition of tiles. If `tile` already has a WCS structure, this function won't do anything.

In many cases, tiles are created for internal/low-level processing. Hence for preformance reasons, when creating the tiles they don't have any WCS structure. When needed, this function can be used to add a WCS structure to each tile tile by copying the WCS structure of its block and correcting the reference point's coordinates within the tile.

`double *` [Function]
`gal_wcs_warp_matrix` (*struct wcsprm* `*wcs`)

Return the Warping matrix of the given WCS structure as an array of double precision floating points. This will be the final matrix, irrespective of the type of storage in the WCS structure. Recall that the FITS standard has several methods to store the matrix. The output is an allocated square matrix with each side equal to the number of dimensions.

`void` [Function]
`gal_wcs_decompose_pc_cdelt` (*struct wcsprm* `*wcs`)

Decompose the `PCi_j` and `CDELTi` elements of `wcs`. According to the FITS standard, in the `PCi_j` WCS formalism, the rotation matrix elements m_{ij} are encoded in the `PCi_j` keywords and the scale factors are encoded in the `CDELTi` keywords. There is also another formalism (the `CDi_j` formalism) which merges the two into one matrix.

However, WCSLIB's internal operations are apparently done in the `PCi_j` formalism. So its outputs are also all in that format by default. When the input is a `CDi_j`, WCSLIB will still read the matrix directly into the `PCi_j` matrix and the `CDELTi` values are set to `1` (one). This function is designed to correct such issues: after it is finished, the `CDELTi` values in `wcs` will correspond to the pixel scale, and the `PCi_j` will correction show the rotation.

`double` [Function]
`gal_wcs_angular_distance_deg` (*double* `r1`, *double* `d1`, *double* `r2`, *double* `d2`)

Return the angular distance (in degrees) between a point located at (`r1`, `d1`) to (`r2`, `d2`). All input coordinates are in degrees. The distance (along a great circle) on a sphere between two points is calculated with the equation below.

$$\cos(d) = \sin(d_1)\sin(d_2) + \cos(d_1)\cos(d_2)\cos(r_1 - r_2)$$

However, since the the pixel scales are usually very small numbers, this function won't use that direct formula. It will be use the Haversine formula (`https://en.wikipedia.org/wiki/Haversine_formula`) which is better considering floating point errors:

$$\frac{\sin^2(d)}{2} = \sin^2\left(\frac{d_1 - d_2}{2}\right) + \cos(d_1)\cos(d_2)\sin^2\left(\frac{r_1 - r_2}{2}\right)$$

`double *` [Function]
`gal_wcs_pixel_scale_deg` (*struct wcsprm* `*wcs`)
Return the pixel scale for each dimension of `wcs` in degrees. The output is an array of double precision floating point type with one element for each dimension.

`double` [Function]
`gal_wcs_pixel_area_arcsec2` (*struct wcsprm* `*wcs`)
Return the pixel area of `wcs` in arcsecond squared.

`void` [Function]
`gal_wcs_world_to_img` (*struct wcsprm* `*wcs`, *double* `*ra`, *double* `*dec`, *double* `**x`, *double* `**y`, *size_t* `size`)
Convert the arrays of input world coordinates (`ra` and `dec`) into arrays of image coordinates (`x` and `y`). Each is assumed to be a separate one-dimensional array of `size` elements. If `*x==NULL` or `*y==NULL`, then space will be allocated for them by this function, otherwise, it is assumed that they already contain the space necessary to write the values.

If you don't need the input values after this conversion any more, you can pass the pointers of the `ra` and `dec` arrays and the outputs will be written into them. This can help to avoid extra allocations and freeing.

`void` [Function]
`gal_wcs_img_to_world` (*struct wcsprm* `*wcs`, *double* `*x`, *double* `*y`, *double* `**ra`, *double* `**dec`, *size_t* `size`)
Convert the arrays of input image coordinates (`x` and `y`) into arrays of world coordinates (`ra` and `dec`). Each is assumed to be a separate one-dimensional array of `size` elements. See `gal_wcs_world_to_img` for more.

10.3.10 Text files (`txt.h`)

FITS files are the primary data container in astronomy. FITS indeed as many useful features, but the most universal and portable format for data storage are plain text files. They can be viewed and edited on any text editor or even on the command-line. Therefore the functions in this section are defined to simplify reading from and writing to plain text files.

Gnuastro defines a simple format for metadata of table columns in a plain text file that is discussed in Section 4.5.2 [Gnuastro text table format], page 69. The functions to get information from, read from and write to plain text files also follow those conventions.

`GAL_TXT_LINESTAT_INVALID` [Macro]
`GAL_TXT_LINESTAT_BLANK` [Macro]
`GAL_TXT_LINESTAT_COMMENT` [Macro]
`GAL_TXT_LINESTAT_DATAROW` [Macro]
Status codes for lines in a plain text file that may be returned by `gal_txt_line_stat`). Lines which have a `#` character as their first non-white character are considered to be comments. Lines with nothing but white space characters are considered blank.

int [Function]
gal_txt_line_stat (*char* *line)

Check the contents of line and see if it is a blank, comment, or data line. The returned values are the macros that start with GAL_TXT_LINESTAT.

gal_data_t * [Function]
gal_txt_table_info (*char* *filename, *size_t* *numcols, *size_t* *numrows)

Store the information of each column in filename into an array of data structures with numcols elements (one data structure for each column) see Section 10.3.5.3 [Arrays of datasets], page 252. The total number of rows in the table is also put into the memory that numrows points to.

Note that other than the character strings (column name, units and comments), nothing in the data structure(s) will be allocated by this function for the actual data (e.g., the array or dsize elements). This function is just for column information. This is a low-level function particular to reading tables in FITS format. It is recommended to use gal_table_info which will allow getting information from a variety of table formats (see Section 10.3.11 [Table input output (table.h)], page 281).

gal_data_t * [Function]
gal_txt_table_read (*char* *filename, *size_t* numrows, *gal_data_t* *colinfo, *gal_list_sizet_t* *indexll, *size_t* minmapsize)

Read the columns given in the list indexll from a plain text table into a linked list of data structures, see Section 10.3.7.3 [List of size_t], page 259, and Section 10.3.7.9 [List of gal_data_t], page 267. If the necessary space for each column is larger than minmapsize, don't keep it in the RAM, but in a file on the HDD/SSD, see the description under the same name in Section 10.3.5.1 [Generic data container (gal_data_t)], page 245.

Note that this is a low-level function, so the output data list is the inverse of the input indexs linked list. It is recommended to use gal_table_read for generic reading of tables in any format, see Section 10.3.11 [Table input output (table.h)], page 281.

gal_data_t * [Function]
gal_txt_image_read (*char* *filename, *size_t* minmapsize)

Read the 2D plain text dataset in filename into a dataset and return the dataset. If the necessary space for the image is larger than minmapsize, don't keep it in the RAM, but in a file on the HDD/SSD, see the description under the same name in Section 10.3.5.1 [Generic data container (gal_data_t)], page 245.

void [Function]
gal_txt_write (*gal_data_t* *cols, *gal_list_str_t* *comment, *char* *filename, *int* dontdelete)

Write cols in a plain text file filename. cols may have one or two dimensions which determines the output:

1D cols is treated as a column and a list of datasets (see Section 10.3.7.9 [List of gal_data_t], page 267): every node in the list is written as one column in a table.

2D `cols` is a two dimensional array, it cannot be treated as a list (only one 2D array can currently be written to a text file). So if `cols->next!=NULL` the next nodes in the list are ignored and will not be written.

If `filename` already exists and `dontdelete` is non-zero, then this function will abort with an error and will not write over the existing file. If `comments!=NULL`, a `#` will be put at the start of each node of the list of strings and will be written in the file before the column meta-data in `filename` (see Section 10.3.7.1 [List of strings], page 257).

Note that this is a low-level function for tables. It is recommended to use `gal_table_write` for generic writing of tables in a variety of formats, see Section 10.3.11 [Table input output (`table.h`)], page 281.

10.3.11 Table input output (`table.h`)

Tables are a collection of one dimensional datasets that are packed together into one file. They are the single most common format to store high-level (processed) information, hence they play a very important role in Gnuastro. For a more thorough introduction, please see Gnuastro's Section 5.3 [Table], page 94. Gnuastro's Table program, and all the other programs that can read from and write into tables, use the functions of this section. For a simple demonstration of using the constructs introduced here, see Section 10.4.4 [Library demo - reading and writing table columns], page 317.

Currently only plain text (see Section 4.5.2 [Gnuastro text table format], page 69) and FITS tables are supported by Gnuastro. However, the low-level table infra-structure is written such that for other formats are also possible and in future releases more formats will be supported, please let us know your priorities so they get higher priorities.

`GAL_TABLE_DEF_WIDTH_STR` [Macro]
`GAL_TABLE_DEF_WIDTH_INT` [Macro]
`GAL_TABLE_DEF_WIDTH_LINT` [Macro]
`GAL_TABLE_DEF_WIDTH_FLT` [Macro]
`GAL_TABLE_DEF_WIDTH_DBL` [Macro]
`GAL_TABLE_DEF_PRECISION_INT` [Macro]
`GAL_TABLE_DEF_PRECISION_FLT` [Macro]
`GAL_TABLE_DEF_PRECISION_DBL` [Macro]

The default width and precision for generic types to use in writing numeric types into a text file (plain text and FITS ASCII tables). When the dataset doesn't have any pre-set width and precision (see `disp_width` and `disp_precision` in Section 10.3.5.1 [Generic data container (`gal_data_t`)], page 245) these will be directly used in C's `printf` command to write the number as a string.

`GAL_TABLE_DISPLAY_FMT_STRING` [Macro]
`GAL_TABLE_DISPLAY_FMT_DECIMAL` [Macro]
`GAL_TABLE_DISPLAY_FMT_UDECIMAL` [Macro]
`GAL_TABLE_DISPLAY_FMT_OCTAL` [Macro]
`GAL_TABLE_DISPLAY_FMT_HEX` [Macro]
`GAL_TABLE_DISPLAY_FMT_FLOAT` [Macro]
`GAL_TABLE_DISPLAY_FMT_EXP` [Macro]

`GAL_TABLE_DISPLAY_FMT_GENERAL` [Macro]

The display format used in C's `printf` to display data of different types. The `_STRING` and `_DECIMAL` are unique for printing strings and signed integers, they are mainly here for completeness. However, unsigned integers and flating points can be displayed in multiple formats:

Unsigned integer
: For unsigned integers, it is possible to choose from `_UDECIMAL` (unsigned decimal), `_OCTAL` (octal notation, for example `125` in decimal will be displayed as `175`), and `_HEX` (hexadecimal notation, for example `125` in decimal will be displayed as `7D`).

Floating point
: For floating point, it is possible to display the number in `_FLOAT` (floating point, for example `1500.345`), `_EXP` (exponential, for example `1.500345e+03`), or `_GENERAL` which is the best of the two for the given number.

`GAL_TABLE_FORMAT_INVALID` [Macro]
`GAL_TABLE_FORMAT_TXT` [Macro]
`GAL_TABLE_FORMAT_AFITS` [Macro]
`GAL_TABLE_FORMAT_BFITS` [Macro]

All the current acceptable table formats to Gnuastro. The `AFITS` and `BFITS` represent FITS ASCII tables and FITS Binary tables. You can use these anywhere you see the `tableformat` variable.

`GAL_TABLE_SEARCH_INVALID` [Macro]
`GAL_TABLE_SEARCH_NAME` [Macro]
`GAL_TABLE_SEARCH_UNIT` [Macro]
`GAL_TABLE_SEARCH_COMMENT` [Macro]

When the desired column is not a number, these values determine if the string to match, or regular expression to search, be in the *name*, *units* or *comments* of the column meta data. These values should be used for the `searchin` variables of the functions.

`gal_data_t *` [Function]
`gal_table_info` (*char* `*filename`, *char* `*hdu`, *size_t* `*numcols`, *size_t* `*numrows`, *int* `*tableformat`)

Store the information of each column in a table (either as a text file or as a FITS table) into an array of data structures with `numcols` structures (one data structure for each column). The number of rows is stored in the space that `numrows` points to. The format of the table (e.g., ascii text file, or FITS binary or ASCII table) will be put in `tableformat` (macros defined above). If the filename is not a FITS file, then `hdu` will not be used (can be `NULL`).

Note that other than the character strings (column name, units and comments), nothing in the data structure(s) will be allocated by this function for the actual data (e.g., the 'array' or 'dsize' elements). This function is just for column information (meta-data), not column contents.

void [Function]
gal_table_print_info (*gal_data_t* *allcols, *size_t* numcols, *size_t* numrows)

This program will print the column information for all the columns (output of gal_table_info). The output is in the same format as this command with Gnuastro Table program (see Section 5.3 [Table], page 94):

```
$ asttable --info table.fits
```

gal_data_t * [Function]
gal_table_read (*char* *filename, *char* *hdu, *gal_list_str_t* *cols, *int* searchin, *int* ignorecase, *int* minmapsize)

Read the specified columns in a text file (named filename) into a linked list of data structures. If the file is FITS, then hdu will also be used, otherwise, hdu is ignored.

The information to search for columns should be specified by the cols list of strings (see Section 10.3.7.1 [List of strings], page 257). The string in each node of the list may be a number, an exact match to a column name, or a regular expression (in GNU AWK format) enclosed in / /. The searchin value must be one of the macros defined above. If cols is NULL, then this function will read the full table.

The output is an individually allocated list of datasets (see Section 10.3.7.9 [List of gal_data_t], page 267) with the same order of the cols list. Note that one column node in the cols list might give multiple columns (for example from regular expressions), in this case, the order of output columns that correspond to that one input, are in order of the table (which column was read first). So the first requested column is the first popped data structure and so on.

void [Function]
gal_table_comments_add_intro (*gal_list_str_t* **comments, *char* *program_string, *time_t* *rawtime)

Add some basic information to the list of comments. This basic information includes the following information

- If the program is run in a Git version controlled directory, Git's description is printed (see description under COMMIT in Section 4.9 [Output headers], page 78).
- The calendar time that is stored in rawtime (time_t is C's calendar time format defined in time.h). You can calculate the time in this format with the following expressions:

```
time_t rawtime;
time(&rawtime);
```

- The name of your program in program_string. If it is NULL, this line is ignored.

void [Function]
gal_table_write (*gal_data_t* *cols, *gal_list_str_t* *comments, *int* tableformat, *char* *filename, *int* dontdelete)

Write the cols list of datasets into a table in filename (see Section 10.3.7.9 [List of gal_data_t], page 267). The format of the table can be determined with tableformat that accepts the macros defined above. If comments!=NULL, then the list of comments will also be printed into the output table. When the output table is a plain text file, each node's string will be printed after a # (so it can be considered

as a comment). If filename already exists and dontdelete is not zero, this function will abort the program with an error.

void [Function]
gal_table_write_log (*gal_data_t* *logll, *char* *program_string, *time_t* *rawtime, *gal_list_str_t* *comments, *char* *filename, *int* dontdelete, *int* quiet)

Write the logll list of datasets into a table in filename (see Section 10.3.7.9 [List of gal_data_t], page 267). This function is just a wrapper around gal_table_comments_add_intro and gal_table_write (see above). If quiet is non-zero, this function will print a message saying that the filename has been created.

10.3.12 Arithmetic on datasets (arithmetic.h)

When the dataset's type and other information are already known, any programming language (including C) provides some very good tools for various operations (including arithmetic operations like addition) on the dataset with a simple loop. However, as an author of a program, making assumptions about the type of data, its dimensionality and other basic characteristics will come with a large processing burdern.

For example if you always read your data as double precision floating points for a simple operation like addition with an integer constant, you will be wasting a lot of CPU and memory when the input dataset is int32 type for example (see Section 4.4 [Numeric data types], page 64). This overhead may be small for small images, but as you scale your process up and work with hundred/thousands of files that can be very large, this overhead will take a significant portion of the processing power. The functions and macros in this section are designed precisely for this purpose: to allow you to do any of the defined operations on any dataset with no overhead (in the native type of the dataset).

Gnuastro's Arithmetic program uses the functions and macros of this section, so please also have a look at the Section 6.2 [Arithmetic], page 108, program and in particular Section 6.2.2 [Arithmetic operators], page 109, for a better description of the operators discussed here.

The main function of this library is gal_arithmetic that is described below. It can take an arbitrary number of arguments as operands (depending on the operator, similar to printf). Its first two arguments are integers specifying the flags and operator. So first we will review the constants for the recognized flags and operators and discuss them, then introduce the actual function.

GAL_ARITHMETIC_INPLACE [Macro]
GAL_ARITHMETIC_FREE [Macro]
GAL_ARITHMETIC_NUMOK [Macro]
GAL_ARITHMETIC_FLAGS_ALL [Macro]

Bit-wise flags to pass onto gal_arithmetic (see below). To pass multiple flags, use the bitwise-or operator, for example GAL_ARITHMETIC_INPLACE | GAL_ARITHMETIC_FREE. GAL_ARITHMETIC_FLAGS_ALL is a combination of all flags to shorten your code if you want all flags activated. Each flag is described below:

GAL_ARITHMETIC_INPLACE

Do the operation in-place (in the input dataset, thus modifying it) to improve CPU and memory usage. If this flag is used, after gal_arithmetic

finishes, the input dataset will be modified. It is thus useful if you have no more need for the input after the operation.

`GAL_ARITHMETIC_FREE`
: Free (all the) input dataset(s) after the operation is done. Hence the inputs are no longer usable after `gal_arithmetic`.

`GAL_ARITHMETIC_NUMOK`
: It is acceptable to use a number and an array together. For example if you want to add all the pixels in an image with a single number you can pass this flag to avoid having to allocate a constant array the size of the image (with all the pixels having the same number).

`GAL_ARITHMETIC_OP_PLUS` [Macro]
`GAL_ARITHMETIC_OP_MINUS` [Macro]
`GAL_ARITHMETIC_OP_MULTIPLY` [Macro]
`GAL_ARITHMETIC_OP_DIVIDE` [Macro]
`GAL_ARITHMETIC_OP_LT` [Macro]
`GAL_ARITHMETIC_OP_LE` [Macro]
`GAL_ARITHMETIC_OP_GT` [Macro]
`GAL_ARITHMETIC_OP_GE` [Macro]
`GAL_ARITHMETIC_OP_EQ` [Macro]
`GAL_ARITHMETIC_OP_NE` [Macro]
`GAL_ARITHMETIC_OP_AND` [Macro]
`GAL_ARITHMETIC_OP_OR` [Macro]

Binary operators (requiring two operands) that accept datasets of any recognized type (see Section 4.4 [Numeric data types], page 64). When `gal_arithmetic` is called with any of these operators, it expects two datasets as arguments. For a full description of these operators with the same name, see Section 6.2.2 [Arithmetic operators], page 109. The first dataset/operand will be put on the left of the operator and the second will be put on the right. The output type of the first four is determined from the input types (largest type of the inputs). The rest (which are all conditional operators) will output a binary `uint8_t` (or `unsigned char`) dataset with values of either `0` (zero) or `1` (one).

`GAL_ARITHMETIC_OP_NOT` [Macro]

The logical NOT operator. When `gal_arithmetic` is called with this operator, it only expects one operand (dataset), since this is a unary operator. The output is `uint8_t` (or `unsigned char`) dataset of the same size as the input. Any non-zero element in the input will be `0` (zero) in the output and any `0` (zero) will have a value of `1` (one).

`GAL_ARITHMETIC_OP_ISBLANK` [Macro]

A unary operator with output that is `1` for any element in the input that is blank, and `0` for any non-blank element. When `gal_arithmetic` is called with this operator, it will only expect one input dataset. The output dataset will have `uint8_t` (or `unsigned char`) type.

`gal_arithmetic` with this operator is just a wrapper for the `gal_blank_flag` function of Section 10.3.4 [Library blank values (`blank.h`)], page 243, and this operator is

just included for completeness in arithmetic operations. So in your program, it might be easier to just call `gal_blank_flag`.

`GAL_ARITHMETIC_OP_WHERE` [Macro]

The three-operand *where* operator thoroughly discussed in Section 6.2.2 [Arithmetic operators], page 109. When `gal_arithmetic` is called with this operator, it will only expect three input datasets: the first (which is the same as the returned dataset) is the array that will be modified. The second is the condition dataset (that must have a `uint8_t` or `unsigned char` type), and the third is the value to be used if condition is non-zero.

As a result, note that the order of operands when calling `gal_arithmetic` with `GAL_ARITHMETIC_OP_WHERE` is the opposite of running Gnuastro's Arithmetic program with the `where` operator (see Section 6.2 [Arithmetic], page 108). This is because the latter uses the reverse-Polish notation which isn't necessary when calling a function (see Section 6.2.1 [Reverse polish notation], page 108).

`GAL_ARITHMETIC_OP_SQRT` [Macro]
`GAL_ARITHMETIC_OP_LOG` [Macro]
`GAL_ARITHMETIC_OP_LOG10` [Macro]

Unary operator functions for calculating the square root ($\sqrt{i}$), $ln(i)$ and $log(i)$ mathematic operators on each element of the input dataset. The output will have the same type as the input, so if your inputs are integer types be careful.

If you want your output to be floating point but your input is an integer type, you can convert the input to a floating point type with `gal_data_copy_to_new_type` or `gal_data_copy_to_new_type_free`(see Section 10.3.5.4 [Copying datasets], page 252).

`GAL_ARITHMETIC_OP_MINVAL` [Macro]
`GAL_ARITHMETIC_OP_MAXVAL` [Macro]
`GAL_ARITHMETIC_OP_NUMVAL` [Macro]
`GAL_ARITHMETIC_OP_SUMVAL` [Macro]
`GAL_ARITHMETIC_OP_MEANVAL` [Macro]
`GAL_ARITHMETIC_OP_STDVAL` [Macro]
`GAL_ARITHMETIC_OP_MEDIANVAL` [Macro]

Unary operand statistical operators that will return a single value for datasets of any size. These are just wrappers around similar functions in Section 10.3.18 [Statistical operations (`statistics.h`)], page 302, and are included in `gal_arithmetic` only for completeness (to use easily in Section 6.2 [Arithmetic], page 108). In your programs, it will probably be easier if you use those `gal_statistics_` functions directly.

`GAL_ARITHMETIC_OP_ABS` [Macro]

Unary operand absolute-value operator.

`GAL_ARITHMETIC_OP_MIN` [Macro]
`GAL_ARITHMETIC_OP_MAX` [Macro]
`GAL_ARITHMETIC_OP_NUM` [Macro]
`GAL_ARITHMETIC_OP_SUM` [Macro]
`GAL_ARITHMETIC_OP_MEAN` [Macro]
`GAL_ARITHMETIC_OP_STD` [Macro]

`GAL_ARITHMETIC_OP_MEDIAN` [Macro]

Multi-operand statistical operations. When `gal_arithmetic` is called with any of these operators, it will expect only a single operand that will be interpretted as a list of datasets (see Section 10.3.7.9 [List of `gal_data_t`], page 267. The output will be a single dataset with each of its elements replaced by the respective statistical operation on the whole list. See the discussion under the `min` operator in Section 6.2.2 [Arithmetic operators], page 109.

`GAL_ARITHMETIC_OP_POW` [Macro]

Binary operator to-power operator. When `gal_arithmetic` is called with any of these operators, it will expect two operands: rasing the first by the second. This operator only accepts floating point inputs and the output is also floating point.

`GAL_ARITHMETIC_OP_BITAND` [Macro]
`GAL_ARITHMETIC_OP_BITOR` [Macro]
`GAL_ARITHMETIC_OP_BITXOR` [Macro]
`GAL_ARITHMETIC_OP_BITLSH` [Macro]
`GAL_ARITHMETIC_OP_BITRSH` [Macro]
`GAL_ARITHMETIC_OP_MODULO` [Macro]

Binary integer-only operand operators. These operators are only defined on integer data types. When `gal_arithmetic` is called with any of these operators, it will expect two operands: the first is put on the left of the operator and the second on the right. The ones starting with `BIT` are the respective bit-wise operators in C and `MODULO` is the modulo/remainder operator. For a discussion on these operators, please see Section 6.2.2 [Arithmetic operators], page 109.

The output type is determined from the input types and C's internal conversions: it is strongly recommended that both inputs have the same type (any integer type), otherwise the bit-wise behavior will be determined by your compiler.

`GAL_ARITHMETIC_OP_BITNOT` [Macro]

The unary bit-wise NOT operator. When `gal_arithmetic` is called with any of these operators, it will expect one operand of an integer type and preform the bitwise-NOT operation on it. The output will have the same type as the input.

`GAL_ARITHMETIC_OP_TO_UINT8` [Macro]
`GAL_ARITHMETIC_OP_TO_INT8` [Macro]
`GAL_ARITHMETIC_OP_TO_UINT16` [Macro]
`GAL_ARITHMETIC_OP_TO_INT16` [Macro]
`GAL_ARITHMETIC_OP_TO_UINT32` [Macro]
`GAL_ARITHMETIC_OP_TO_INT32` [Macro]
`GAL_ARITHMETIC_OP_TO_UINT64` [Macro]
`GAL_ARITHMETIC_OP_TO_INT64` [Macro]
`GAL_ARITHMETIC_OP_TO_FLOAT32` [Macro]
`GAL_ARITHMETIC_OP_TO_FLOAT64` [Macro]

Unary type-conversion operators. When `gal_arithmetic` is called with any of these operators, it will expect one operand and convert it to the requested type. Note that with these operators, `gal_arithmetic` is just a wrapper over the `gal_data_copy_to_new_type` or `gal_data_copy_to_new_type_free` that are

discussed in `Copying datasets`. It accepts these operators only for completeness and easy usage in Section 6.2 [Arithmetic], page 108. So in your programs, it might be preferrable to directly use those functions.

`gal_data_t *` [Function]
`gal_arithmetic` (*int operator, unsigned char flags, ...*)

Do the arithmetic operation of `operator` on the given operands (the third argument and any further argument). Certain special conditions can also be specified with the `flag` operator. The acceptable values for `operator` are defined in the macros above.

`gal_arithmetic` is a multi-argument function (like C's `printf`). In other words, the number of necessary arguments is not fixed and depends on the value to `operator`. Here are a few examples showing this variability:

```
out_1=gal_arithmetic(GAL_ARITHMETIC_OP_LOG,   0, in_1);
out_2=gal_arithmetic(GAL_ARITHMETIC_OP_PLUS,  0, in_1, in_2);
out_3=gal_arithmetic(GAL_ARITHMETIC_OP_WHERE, 0, in_1, in_2, in_3);
```

The number of necessary operands for each operator (and thus the number of necessary arguments to `gal_arithmetic`) are described above under each operator.

10.3.13 Tessellation library (`tile.h`)

In many contexts, it is desirable to slice the dataset into subsets or tiles (overlapping or not). In such a way that you can work on each tile independently. One method would be to copy that region to a separate allocated space, but in many contexts this isn't necessary and infact can be a big burden on CPU/Memory usage. The `block` pointer in Gnuastro's Section 10.3.5.1 [Generic data container (`gal_data_t`)], page 245, is defined for such situations: where allocation is not necessary. You just want to read the data or write to it independently (or in coordination with) other regions of the dataset. Added with parallel processing, this can greatly improve the time/memory consumption.

See the figure below for example: assume the `larger` dataset is a contiguous block of memory that you are interpretting as a 2D array. But you only want to work on the smaller `tile` region.

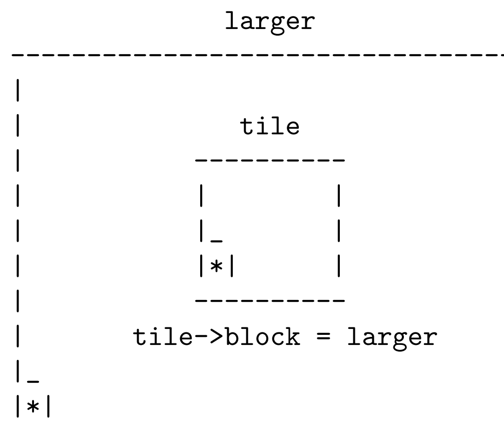

To use `gal_data_t`'s `block` concept, you allocate a `gal_data_t *tile` which is initialized with the pointer to the first element in the sub-array (as its `array` argument). Note that this is not necessarily the first element in the larger array. You can set the size of the

tile along with the initialization as you please. Recall that, when given a non-NULL pointer as array, gal_data_initialize (and thus gal_data_alloc) do not allocate any space and just uses the given pointer for the new array element of the gal_data_t. So your tile data structure will not be pointing to a separately allocated space.

After the allocation is done, you just point tile->block to the larger dataset which hosts the full block of memory. Where relevant, Gnuastro's library functions will check the block pointer of their input dataset to see how to deal with dimensions and increments so they can always remain within the tile. The tools introduced in this section are designed to help in defining and working with tiles that are created in this manner.

Since the block structure is defined as a pointer, arbitrary levels of tesselation/griding are possible (tile->block may itself be a tile in an even larger allocated space). Therefore, just like a linked-list (see Section 10.3.7 [Linked lists (list.h)], page 255), it is important to have the block pointer of the largest (allocated) dataset set to NULL. Normally, you won't have to worry about this, because gal_data_initialize (and thus gal_data_alloc) will set the block element to NULL by default, just remember not to change it. You can then only change the block element for the tiles you define over the allocated space.

Below, we will first review constructs for Section 10.3.13.1 [Independent tiles], page 289, and then define the current approach to fully tessellating a dataset (or covering every pixel/data-element with a non-overlapping tile grid in Section 10.3.13.2 [Tile grid], page 294. This approach to dealing with parts of a larger block was inspired from a similarly named concept in the GNU Scientific Library (GSL), see its "Vectors and Matrices" chapter for their implementation.

10.3.13.1 Independent tiles

The most general application of tiles is to treat each independently, for example they may overlap, or they may not cover the full image. This section provides functions to help in checking/inspecting such tiles. In Section 10.3.13.2 [Tile grid], page 294, we will discuss functions that define/work-with a tile grid (where the tiles don't overlap and fully cover the input dataset). Therefore, the functions in this section are general and can be used for the tiles produced by that section also.

void [Function]
gal_tile_start_coord (*gal_data_t* *tile, *size_t* *start_coord)

> Calculate the starting coordinates of a tile in its allocated block of memory and write them in the memory that start_coord points to (which must have tile->ndim elements).

void [Function]
gal_tile_start_end_coord (*gal_data_t* *tile, *size_t* *start_end, *int* rel_block)

> Put the starting and ending (end point is not inclusive) coordinates of tile into the start_end array. It is assumed that a space of 2*tile->ndim has been already allocated (static or dynamic) for start_end before this function is called.
>
> rel_block (or relative-to-block) is only relevant when tile has an intermediate tile between it and the allocated space (like a channel, see gal_tile_full_two_layers). If it doesn't (tile->block points the allocated dataset), then the value to rel_block is irrelevant.

When `tile->block` is its self a larger block and `rel_block` is set to 0, then the starting and ending positions will be based on the position within `tile->block`, not the allocated space.

`void *` [Function]
`gal_tile_start_end_ind_inclusive` (*gal_data_t* `*tile`, *gal_data_t* `*work`, *size_t* `*start_end_inc`)

Put the indexs of the first/start and last/end pixels (inclusive) in a tile into the `start_end` array (that must have two elements). NOTE: this function stores the index of each point, not its coordinates. It will then return the pointer to the start of the tile in the `work` data structure (which doesn't have to be equal to `tile->block`.

The outputs of this function are defined to make it easy to parse over an n-dimensional tile. For example, this function is one of the most important parts of the internal processing of in `GAL_TILE_PARSE_OPERATE` function-like macro that is described below.

`gal_data_t *` [Function]
`gal_tile_series_from_minmax` (*gal_data_t* `*block`, *size_t* `*minmax`, *size_t* `number`)

Construct a list of tile(s) given coordinates of the minimum and maximum of each tile. The minimum and maximums are assumed to be inclusive. The returned pointer is an allocated `gal_data_t` array that can later be freed with `gal_data_array_free` (see Section 10.3.5.3 [Arrays of datasets], page 252). Internally, each element of the output array points to the next element, so the output may also be treated as a list of datasets (see Section 10.3.7.9 [List of `gal_data_t`], page 267) and passed onto the other functions described in this section.

The array keeping the minmium and maximum coordinates for each tile must have the following format. So in total `minmax` must have `2*ndim*number` elements.

```
| min0_d0 | min0_d1 | max0_d0 | max0_d1 | ...
                  ... | minN_d0 | minN_d1 | maxN_d0 | maxN_d1 |
```

`gal_data_t *` [Function]
`gal_tile_block` (*gal_data_t* `*tile`)

Return the dataset that contains `tile`'s allocated block of memory. If tile is immediately defined as part of the allocated block, then this is equivalent to `tile->block`. However, it is possible to have multiple layers of tiles (where `tile->block` is itself a tile). So this function is the most generic way to get to the actual allocated dataset.

`size_t` [Function]
`gal_tile_block_increment` (*gal_data_t* `*block`, *size_t* `*tsize`, *size_t* `num_increment`, *size_t* `*coord`)

Return the increment necessary to start at the next contiguous patch memory associated with a tile. `block` is the allocated block of memory and `tsize` is the size of the tile along every dimension. If `coord` is `NULL`, it is ignored. Otherwise, it will contain the coordinate of the start of the next contiguous patch of memory.

This function is intended to be used in a loop and `num_increment` is the main variable to this function. For the first time you call this function, it should be `1`. In subsequent calls (while you are parsing a tile), it should be increased by one.

gal_data_t * [Function]
gal_tile_block_write_const_value (*gal_data_t* ***tilevalues**, *gal_data_t* ***tilesll**, *int* **initialize**)

Write a constant value for each tile over the area it covers in an allocated dataset that is the size of `tile`'s allocated block of memory (found through `gal_tile_block` described above). The arguments to this function are:

`tilevalues`
: This must be an array that has the same number of elements as the nodes in in `tilesll` and in the same order that 'tilesll' elements are parsed (from top to bottom, see Section 10.3.7 [Linked lists (`list.h`)], page 255). As a result the dimensionality of this array is irrelevant, it will be parsed contiguously.

`tilesll`
: The list of input tiles (see Section 10.3.7.9 [List of `gal_data_t`], page 267). Internally, it might be stored as an array (for example the output of `gal_tile_series_from_minmax` described above), but this function doesn't care, it will parse the `next` elements to go to the next tile. This function will not pop-from or free the `tilesll`, it will only parse it from start to end.

`initialize`
: Initialize the allocated space with blank values before writing in the constant values. This can be useful when the tiles don't cover the full allocated block.

gal_data_t * [Function]
gal_tile_block_check_tiles (*gal_data_t* ***tilesll**)

Make a copy of the memory block and fill it with the index of each tile in `tilesll` (counting from 0). The non-filled areas will have blank values. The output dataset will have a type of `GAL_TYPE_INT32` (see Section 10.3.3 [Library data types (`type.h`)], page 238).

This function can be used when you want to check the coverage of each tile over the allocated block of memory. It is just a wrapper over the `gal_tile_block_write_const_value`.

void * [Function]
gal_tile_block_relative_to_other (*gal_data_t* ***tile**, *gal_data_t* ***other**)

Return the pointer corresponding to the start of the region covered by `tile` over the `other` dataset. See the examples in `GAL_TILE_PARSE_OPERATE` for some example applications of this function.

void [Function]
gal_tile_block_blank_flag (*gal_data_t* ***tilell**, *size_t* **numthreads**)

Check if each tile in the list has blank values and update its `flag` to mark this check and its result (see Section 10.3.5.1 [Generic data container (`gal_data_t`)], page 245). The operation will be done on `numthreads` threads.

GAL_TILE_PARSE_OPERATE (IN, OTHER, PARSE_OTHER, CHECK_BLANK, OP) [Function-like macro]

Parse over IN (which can be a tile or a fully allocated block of memory) and do the OP operation on it (can be any combination of C expresssions). If OTHER!=NULL, this macro will allow access to its element(s) and it can optionally be parsed while parsing over IN. You can use this to parse the same region over two arrays. The input arguments are (the expected types of each argument are also written here):

IN (gal_data_t)
: Input dataset, this can be a tile or an allocated block of memory.

OTHER (gal_data_t)
: Dataset (gal_data_t) to parse along with IN. It can be NULL. In that case, o (see description of OP below) will be NULL and should not be used. If PARSE_OTHER is zero, the size of this dataset is irrelevant. Otherwise, it has to have the same size as the allocated block of IN.

PARSE_OTHER (int)
: Parse the other dataset along with the input. When this is non-zero and OTHER!=NULL, then the o pointer will be incremented to cover the OTHER tile at the same rate as i, see description of OP for i and o.

CHECK_BLANK (int)
: If it is non-zero, then the input will be checked for blank values and OP will only be called when we are not on a blank element.

OP
: Operator: this can be any number of C experssions. This macro is going to define a itype *i variable which will increment over each element of the input array/tile. itype will be replaced with the C type that corresponds to the type of INPUT. As an example, if INPUT's type is GAL_DATA_UINT16 or GAL_DATA_FLOAT32, i will be defined as uint16 or float respectively.

  This function-like macro will also define an otype *o which you can use to access an element of the OTHER dataset (if OTHER!=NULL). o will correspond to the type of OTHER (similar to itype and INPUT discussed above). If PARSE_OTHER is non-zero, then o will also be incremented to the same index element but in the other array. You can use these along with any other variable you define before this macro to process the input and store it in the other.

  All variables within this function-like macro begin with tpo_ except for the three variables listed above. So as long as you don't start the names of your variables with this prefix everything will be fine. Note that i (and possibly o) will be incremented once by this function-like macro, so don't increment them within OP. As a summary, the three variables you have access to within OP are:

  i
  : Pointer to the element of INPUT that is being parsed with the proper type.

`o` Pointer to the element of `OTHER` that is being parsed with the proper type. `o` can only be used if `OTHER!=NULL` and it will be parsed/incremented if `PARSE_OTHER` is non-zero.

`b` Blank value in the type of `INPUT`.

You can use a given tile (`tile` on a dataset that it was not initialized with (but has the same size, let's call it `new`) with the following steps:

```
void *tarray;
gal_data_t *tblock;

/* 'tile->block' must corrected AFTER 'tile->array' */
tarray      = tile->array;
tblock      = tile->block;
tile->array = gal_tile_block_relative_to_other(tile, new);
tile->block = new;

/* Parse and operate over this region of the 'new' dataset. */
GAL_TILE_PARSE_OPERATE(tile, NULL, 0, 0, {
    YOUR_PROCESSING;
  });

/* Reset 'tile->block' and 'tile->array'. */
tile->array=tarray;
tile->block=tblock;
```

You can work on the same region of another block in one run of this function-like macro. To do that, you can make a fake tile and pass that as the `OTHER` argument. Below is a demonstration, `tile` is the actual tile that you start with and `new` is the other block of allocated memory.

```
size_t zero=0;
gal_data_t *faketile;

/* Allocate the fake tile, these can be done outside a loop
 * (over many tiles). */
faketile=gal_data_alloc(NULL, new->type, 1, &zero,
                        NULL, 0, -1, NULL, NULL, NULL);
free(faketile->array);               /* To keep things clean. */
free(faketile->dsize);               /* To keep things clean. */
faketile->block = new;
faketile->ndim  = new->ndim;

/* These can be done in a loop (over many tiles). */
faketile->size  = tile->size;
faketile->dsize = tile->dsize;
faketile->array = gal_tile_block_relative_to_other(tile, new);

// Do your processing....
```

```
GAL_TILE_PARSE_OPERATE(tile, faketile, 1, 1, {
    YOUR_PROCESSING_EXPRESSIONS;
  });

// Clean up.
bintile->array=NULL;
bintile->dsize=NULL;
gal_data_free(bintile);
```

10.3.13.2 Tile grid

One very useful application of tiles is to completely cover an input dataset with tiles. Such that you know every pixel/data-element of the input image is covered by only one tile. The constructs in this section allow easy definition of such a tile structure. They will create lists of tiles that are also usable by the general tools discussed in Section 10.3.13.1 [Independent tiles], page 289.

As discussed in Section 4.6 [Tessellation], page 72, (mainly raw) astronomical images will mostly require two layers of tessellation, one for amplifier channels which all have the same size and another (smaller tile-size) tessellation over each channel. Hence, in this section we define a general structure to keep the main parameters of this two-layer tessellation and help in benefiting from it.

`gal_tile_two_layer_params` [Type (C `struct`)]

This is the general structure to keep all the necessary parameters for a two-layer tessellation.

```
struct gal_tile_two_layer_params
{
  /* Inputs */
  size_t            *tilesize;  /******************************/
  size_t         *numchannels;  /* These parameters have to be */
  float         remainderfrac;  /* filled manually before      */
  uint8_t         workoverch;   /* calling the functions in    */
  uint8_t          checktiles;  /* this section.               */
  uint8_t      oneelempertile;  /******************************/

  /* Internal parameters. */
  size_t                 ndim;
  size_t             tottiles;
  size_t         tottilesinch;
  size_t          totchannels;
  size_t         *channelsize;
  size_t            *numtiles;
  size_t        *numtilesinch;
  char         *tilecheckname;
  size_t         *permutation;
  size_t          *firsttsize;

  /* Tile and channel arrays (which are also lists). */
```

```
    gal_data_t           *tiles;
    gal_data_t        *channels;
  };
```

size_t * [Function]
gal_tile_full (*gal_data_t* *input, *size_t* *regular, *float* remainderfrac, *gal_data_t* **out, *size_t* multiple, *size_t* **firsttsize)

Cover the full dataset with (mostly) identical tiles and return the number of tiles created along each dimension. The regular tile size (along each dimension) is determined from the `regular` array. If `input`'s size is not an exact multiple of `regular` for each dimension, then the tiles touching the edges in that dimension will have a different size to fully cover every element of the input (depending on `remainderfrac`).

The output is an array with the same dimensions as `input` which contains the number of tiles along each dimension. See Section 4.6 [Tessellation], page 72, for a description of its application in Gnuastro's programs and `remainderfrac`, just note that this function defines only onelayer of tiles.

This is a low-level function (independent of the `gal_tile_two_layer_params` strcuture defined above). If you want a two-layer tessellation, directly call `gal_tile_full_two_layers` that is described below. The input arguments to this function are:

`input` The main dataset (allocated block) which you want to create a tessellation over (only used for its sizes). So `input` may be a tile also.

`regular` The the size of the regular tiles along each of the input's dimensions. So it must have the same number of elements as the dimensions of `input` (or `input->ndim`).

`remainderfrac`
The significant fraction of the remainder space to see if it should be split into two and put on both sides of a dimension or not. This is thus only relevant `input` length along a dimension isn't an exact multiple of the regular tile size along that dimension. See Section 4.6 [Tessellation], page 72, for a more thorough discussion.

`out` Pointer to the array of data structures that will keep all the tiles (see Section 10.3.5.3 [Arrays of datasets], page 252). If `*out==NULL`, then the necessary space to keep all the tiles will be allocated. If not, then all the tile information will be filled from the dataset that `*out` points to, see `multiple` for more.

`multiple` When `*out==NULL` (and thus will be allocated by this function), allocate space for `multiple` times the number of tiles needed. This can be very useful when you have several more identically sized `inputs`, and you want all their tiles to be allocated (and thus indexed) together, even though they have different `block` datasets (that then link to one allocated space). See the definition of channels in Section 4.6 [Tessellation], page 72, and `gal_tile_full_two_layers` below.

`firsttsize`
The size of the first tile along every dimension. This is only different from the regular tile size when `regular` is not an exact multiple of `input`'s length along every dimension. This array is allocated internally by this function.

void [Function]
gal_tile_full_sanity_check (*char* `*filename`, *char* `*hdu`, *gal_data_t* `*input`, *struct gal_tile_two_layer_params* `*tl`)

Make sure that the input parameters (in `tl`, short for two-layer) correspond to the input dataset. `filename` and `hdu` are only required for error messages. Also, allocate and fill the `tl->channelsize` array.

void [Function]
gal_tile_full_two_layers (*gal_data_t* `*input`, *struct gal_tile_two_layer_params* `*tl`)

Create the two layered tessellation in `tl`. The general set of steps you need to take to define the two-layered tessellation over an image can be seen in the example code below.

```
gal_data_t *input;
struct gal_tile_two_layer_params tl;
char *filename="input.fits", *hdu="1";

/* Set all the inputs shown in the structure definition. */
...

/* Read the input dataset. */
input=gal_fits_img_read(filename, hdu, -1);

/* Do a sanity check and preparations. */
gal_tile_full_sanity_check(filename, hdu, input, &tl);

/* Build the two-layer tessellation*/
gal_tile_full_two_layers(input, &tl);

/* 'tl.tiles' and 'tl.channels' are now a lists of tiles.*/
```

void [Function]
gal_tile_full_permutation (*struct gal_tile_two_layer_params* `*tl`)

Make a permutation to allow the conversion of tile location in memory to its location in the full input dataset and put it in `tl->permutation`. If a permutation has already been defined for the tessellation, this function will not do anything. If permutation won't be necessary (there is only one channel or one dimension), then this function will not do anything (`tl->permutation` must have been initialized to `NULL`).

When there is only one channel OR one dimension, the tiles are allocated in memory in the same order that they represent the input data. However, to make channel-independent processing possible in a generic way, the tiles of each channel are allocated

contiguously. So, when there is more than one channel AND more than one dimension, the index of the tile does not correspond to its position in the grid covering the input dataset.

The example below may help clarify: assume you have a 6x6 tessellation with two channels in the horizontal and one in the vertical. On the left you can see how the tile IDs correspond to the input dataset. NOTE how '03' is on the second row, not on the first after '02'. On the right, you can see how the tiles are stored in memory (and shown if you simply write the array into a FITS file for example).

```
Corresponding to input              In memory
----------------------            --------------
  15 16 17 33 34 35              30 31 32 33 34 35
  12 13 14 30 31 32              24 25 26 27 28 29
  09 10 11 27 28 29              18 19 20 21 22 23
  06 07 08 24 25 26     <--      12 13 14 15 16 17
  03 04 05 21 22 23              06 07 08 09 10 11
  00 01 02 18 19 20              00 01 02 03 04 05
```

As a result, if your values are stored in same order as the tiles, and you want them in over-all memory (for example to save as a FITS file), you need to permute the values:

```
gal_permutation_apply(values, tl->permutation);
```

If you have values over-all and you want them in tile-order, you can apply the inverse permutation:

```
gal_permutation_apply_inverse(values, tl->permutation);
```

Recall that this is the definition of permutation in this context:

```
permute:    IN_ALL[ i       ]   =   IN_MEMORY[ perm[i] ]
inverse:    IN_ALL[ perm[i] ]   =   IN_MEMORY[ i       ]
```

void [Function]
`gal_tile_full_values_write` (*gal_data_t* `*tilevalues`, *struct gal_tile_two_layer_params* `*tl`, *char* `*filename`, *gal_fits_list_key_t* `*keys`, *char* `*program_string`)

Write one value for each tile into a file. It is important to note that the values in `tilevalues` must be ordered in the same manner as the tiles, so `tilevalues->array[i]` is the value that should be given to `tl->tiles[i]`. The `tl->permutation` array must have been initialized before calling this function with `gal_tile_full_permutation`.

gal_data_t * [Function]
`gal_tile_full_values_smooth` (*gal_data_t* `*tilevalues`, *struct gal_tile_two_layer_params* `*tl`, *size_t* `width`, *size_t* `numthreads`)

Smooth the given values with a flat kernel of the given `width`. This cannot be done manually because if `tl->workoverch==0`, tiles in different channels must not be mixed/smoothed. Also the tiles are contiguous within the channel, not within the image, see the description under `gal_tile_full_permutation`.

size_t [Function]
gal_tile_full_id_from_coord (*struct gal_tile_two_layer_params* *tl, *size_t* *coord)

Return the ID of the tile that corresponds to the coordinates `coord`. Having this ID, you can use the `tl->tiles` array to get to the proper tile or read/write a value into an array that has one value per tile.

void [Function]
gal_tile_full_free_contents (*struct gal_tile_two_layer_params* *tl)

Free all the allocated arrays within `tl`.

10.3.14 Bounding box (box.h)

Functions related to reporting a the bounding box of certain inputs are declared in `gnuastro/box.h`. All coordinates in this header are in the FITS format (first axis is the horizontal and the second axis is vertical).

void [Function]
gal_box_ellipse_in_box (*double* a, *double* b, *double* theta_rad, *long* *width)

Any ellipse can be enclosed into a rectangular box. The purpose of this function is to give the height and width of that box. `a` is the ellipse major axis, `b` is the minor axis, `theta_rad` is the position angle in radians. The `width` array will contain the output size in long integer type. `width[0]`, and `width[1]` are the number of pixels along the first and second FITS axis.

void [Function]
gal_box_border_from_center (*double* xc, *double* yc, *long* *width, *long* *fpixel, *long* *lpixel)

Given the center (`xc` and `yc`) and width (two element array) of a box, return the coordinates of the first `fpixel` and last `lpixel` pixels (both are two element arrays.

int [Function]
gal_box_overlap (*long* *naxes, *long* *fpixel_i, *long* *lpixel_i, *long* *fpixel_o, *long* *lpixel_o)

We have an image of size `naxes` and want to get the overlap of the image with a box. This function will return 1 if there is an overlap and 0 if there isn't. The input and output box are specified by their first and last pixels. When there is an overlap, the coordinates of the first and last pixels of the overlap will be put in `fpixel_o` and `lpixel_o`.

10.3.15 Polygons (polygon.h)

Polygons are commonly necessary in image processing. In Gnuastro, they are used in Crop (see Section 6.1 [Crop], page 97) for cutting out non-rectangular regions of a image. Warp (see Section 6.4 [Warp], page 138) uses them to warp the images into a new pixel grid. The polygon related Gnuastro library macros and functions are introduced here.

In all the functions here the vertices (and points) are defined as an array. So a polygon with 4 vertices will be identified with an array of 8 elements with the first two elements keeping the 2D coordinates of the first vertice and so on.

`GAL_POLYGON_MAX_CORNERS` [Macro]

The largest number of vertices a polygon can have in this library.

`GAL_POLYGON_ROUND_ERR` [Macro]

We have to consider floating point round-off errors when dealing with polygons. For example we will take `A` as the maximum of `A` and `B` when `A>B-GAL_POLYGON_ROUND_ERR`.

`void` [Function]
`gal_polygon_ordered_corners` (*double* `*in`, *size_t* `n`, *size_t* `*ordinds`)

We have a simple polygon (that can result from projection, so its edges don't collide or it doesn't have holes) and we want to order its corners in an anticlockwise fashion. This is necessary for clipping it and finding its area later. The input vertices can have practically any order.

The input (`in`) is an array containing the coordinates (two values) of each vertice. `n` is the number of corners. So `in` should have `2*n` elements. The output (`ordinds`) is an array with `n` elements specifying the indexs in order. This array must have been allocated before calling this function. The indexes are output for more generic usage, for example in a homographic transform (necessary in warping an image, see Section 6.4.1 [Warping basics], page 139), the necessary order of vertices is the same for all the pixels. In other words, only the positions of the vertices change, not the way they need to be ordered. Therefore, this function would only be necessary once.

As a summary, the input is unchanged, only `n` values will be put in the `ordinds` array. Such that calling the input coordinates in the following fashion will give an anti-clockwise order when there are 4 vertices:

```
1st vertice: in[ordinds[0]*2], in[ordinds[0]*2+1]
2nd vertice: in[ordinds[1]*2], in[ordinds[1]*2+1]
3rd vertice: in[ordinds[2]*2], in[ordinds[2]*2+1]
4th vertice: in[ordinds[3]*2], in[ordinds[3]*2+1]
```

The implementation of this is very similar to the Graham scan in finding the Convex Hull. However, in projection we will never have a concave polygon (the left condition below, where this algorithm will get to E before D), we will always have a convex polygon (right case) or E won't exist!

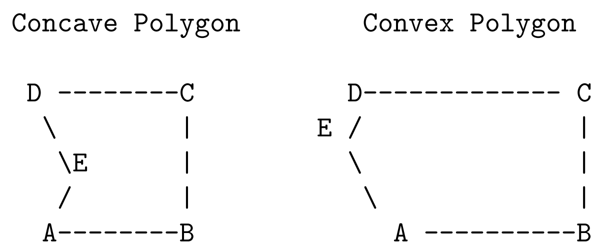

This is because we are always going to be calculating the area of the overlap between a quadrilateral and the pixel grid or the quadrilaterral its self.

The `GAL_POLYGON_MAX_CORNERS` macro is defined so there will be no need to allocate these temporary arrays separately. Since we are dealing with pixels, the polygon can't really have too many vertices.

double [Function]
gal_polygon_area (*double* *v, *size_t* n)
Find the area of a polygon with vertices defined in v. v points to an array of doubles which keep the positions of the vertices such that v[0] and v[1] are the positions of the first vertice to be considered.

int [Function]
gal_polygon_pin (*double* *v, *double* *p, *size_t* n)
Return 1 if the point p is within the polygon whose vertices are defined by v and 0 otherwise. Note that the vertices of the polygon have to be sorted in an anti-clock-wise manner.

int [Function]
gal_polygon_ppropin (*double* *v, *double* *p, *size_t* n)
Similar to gal_polygon_pin, except that if the point p is on one of the edges of a polygon, this will return 0.

void [Function]
gal_polygon_clip (*double* *s, *size_t* n, *double* *c, *size_t* m, *double* *o, *size_t* *numcrn)
Clip (find the overlap of) two polygons. This function uses the Sutherland-Hodgman (https://en.wikipedia.org/wiki/Sutherland%E2%80%93Hodgman_algorithm) polygon clipping algorithm. Note that the vertices of both polygons have to be sorted in an anti-clock-wise manner.

10.3.16 Qsort functions (qsort.h)

The C programming language comes with the qsort (Quick sort) function. qsort is a generic function which allows you to sort any kind of data structure (not just a single number). To to define greater and smaller (for sorting), qsort needs another function, even for simple numerical types. To facilitate numerical sorting for Gnuastro's programs/libraries, Gnuastro defines a function for each type's increasing and decreasing function. You can pass these functions to qsort when your array has the respective type (see Section 4.4 [Numeric data types], page 64).

gal_qsort_index_arr [Global variable]
Pointer to a floating point array (float *) to use as a reference in gal_qsort_index_float_decreasing, see the explanation there for more.

int [Function]
gal_qsort_index_float_decreasing (*const void* *a, *const void* *b)
When passed to qsort, this function will sort a size_t array based on decreasing values in the gal_qsort_index_arr single precision floating point array. The floating point array will not be changed, it is only read. For example, if we have the following source code:

```
#include <stdio.h>
#include <stdlib.h>           /* qsort is defined in stdlib.h. */
#include <gnuastro/qsort.h>
```

```
int
main (void)
{
  size_t s[4]={0, 1, 2, 3};
  float f[4]={1.3,0.2,1.8,0.1};
  gal_qsort_index_arr=f;
  qsort(s, 4, sizeof(size_t), gal_qsort_index_float_decreasing);
  printf("%zu, %zu, %zu, %zu\n", s[0], s[1], s[2], s[3]);
  return EXIT_SUCCESS;
}
```

The output will be: `2, 0, 1, 3`.

`int` [Function]
`gal_qsort_TYPE_increasing` (*const void* `*a`, *const void* `*b`)

When passed to `qsort`, this function will sort an `TYPE` array in increasing order (first element will be the smallest). Please replace `TYPE` (in the function name) with one of the Section 4.4 [Numeric data types], page 64, for example `gal_qsort_int32_increasing`, or `gal_qsort_float64_increasing`.

`int` [Function]
`gal_qsort_TYPE_decreasing` (*const void* `*a`, *const void* `*b`)

When passed to `qsort`, this function will sort an `TYPE` array in decreasing order (first element will be the largest). Please replace `TYPE` (in the function name) with one of the Section 4.4 [Numeric data types], page 64, for example `gal_qsort_int32_decreasing`, or `gal_qsort_float64_decreasing`.

10.3.17 Permutations (`permutation.h`)

Permutation is the technical name for re-ordering of values. The need for permutations occurs a lot during (mainly low-level) processing. To do permutation, you must provide two inputs: an array of values (that you want to re-order inplace) and a permutation array which contains the new index of each element (let's call it `perm`). The diagram below shows the input array before and after the re-ordering.

```
permute:    AFTER[ i       ] = BEFORE[ perm[i] ]     i = 0 .. N-1
inverse:    AFTER[ perm[i] ] = BEFORE[ i       ]     i = 0 .. N-1
```

The functions here are a re-implementation of the GNU Scientific Library's `gsl_permute` function. The reason we didn't use that function was that it uses system-specific types (like `long` and `int`) which can have different widths on different systems, hence are not easily convertable to Gnuastro's fixed width types (see Section 4.4 [Numeric data types], page 64). There is also a separate function for each type, heavily using macros to allow a `base` function to work on all the types. Thus it is hard to read/understand. Hence, Gnuastro contains a re-write of their steps in a new type-agnostic method which is a single function that can work on any type.

As described in GSL's source code and manual, this implementation comes from Donald Knuth's *Art of computer programming* book, in the "Sorting and Searching" chapter of Volume 3 (3rd ed). Exercise 10 of Section 5.2 defines the problem and in the answers, Knuth describes the solution. So if you are interested, please have a look there for more.

We are in contact with the GSL developers and in the future[20] we will submit these implementations to GSL. If they are finally incorporated there, we will delete this section in future versions.

void [Function]
`gal_permutation_check` (*size_t* `*permutation`, *size_t* `size`)

Print how `permutation` will re-order an array that has `size` elements for each element in one one line.

void [Function]
`gal_permutation_apply` (*gal_data_t* `*input`, *size_t* `*permutation`)

Apply `permutation` on the `input` dataset (can have any type), see above for the definition of permutation.

void [Function]
`gal_permutation_apply_inverse` (*gal_data_t* `*input`, *size_t* `*permutation`)

Apply the inverse of `permutation` on the `input` dataset (can have any type), see above for the definition of permutation.

10.3.18 Statistical operations (`statistics.h`)

After reading a dataset into memory from a file or fully simulating it with another process, the most common processes that will be done on it are statistical operations to let you quantify different aspects of the data. the functions in this section describe Gnuastro's current set of tools for this job. All these functions can work on any numeric data type natively (see Section 4.4 [Numeric data types], page 64) and can also work on tiles over a dataset. Hence the inputs and outputs are in Gnuastro's Section 10.3.5.1 [Generic data container (`gal_data_t`)], page 245.

`GAL_STATISTICS_SIG_CLIP_MAX_CONVERGE` [Macro]

The maximum number of clips, when σ-clipping should be done by convergence. If the clipping does not converge before making this many clips, all sigma-clipping outputs will be NaN.

`GAL_STATISTICS_MODE_GOOD_SYM` [Macro]

The minimum acceptable symmetricity of the mode calculation. If the symmetricity of the derived mode is less than this value, all the returned values by `gal_statistics_mode` will have a value of NaN.

`GAL_STATISTICS_SORTED_NOT` [Macro]
`GAL_STATISTICS_SORTED_INCREASING` [Macro]
`GAL_STATISTICS_SORTED_DECREASING` [Macro]

Macros used to identify if the dataset is sorted and increasing, sorted and decreasing or not sorted.

`GAL_STATISTICS_BINS_INVALID` [Macro]
`GAL_STATISTICS_BINS_REGULAR` [Macro]

[20] Gnuastro's Task 14497 (`http://savannah.gnu.org/task/?14497`). If this task is still "postponed" when you are reading this and you are interested to help, your help would be very welcome. Both Gnuastro and GSL developers are very busy, hence both would appreciate your help.

`GAL_STATISTICS_BINS_IRREGULAR` [Macro]

Macros used to identify if the regularity of the bins when defining bins.

`gal_data_t *` [Function]
`gal_statistics_number` (*gal_data_t* `*input`)

Return a single-element `uint64` dataset containing the number of non-blank elements in `input`.

`gal_data_t *` [Function]
`gal_statistics_minimum` (*gal_data_t* `*input`)

Return a single-element dataset containing the minimum non-blank value in `input`. The numerical datatype of the output is the same as `input`.

`gal_data_t *` [Function]
`gal_statistics_maximum` (*gal_data_t* `*input`)

Return a single-element dataset containing the maximum non-blank value in `input`. The numerical datatype of the output is the same as `input`.

`gal_data_t *` [Function]
`gal_statistics_sum` (*gal_data_t* `*input`)

Return a single-element (`double` or `float64`) dataset containing the sum of the non-blank values in `input`.

`gal_data_t *` [Function]
`gal_statistics_mean` (*gal_data_t* `*input`)

Return a single-element (`double` or `float64`) dataset containing the mean of the non-blank values in `input`.

`gal_data_t *` [Function]
`gal_statistics_std` (*gal_data_t* `*input`)

Return a single-element (`double` or `float64`) dataset containing the standard deviation of the non-blank values in `input`.

`gal_data_t *` [Function]
`gal_statistics_mean_std` (*gal_data_t* `*input`)

Return a two-element (`double` or `float64`) dataset containing the mean and standard deviation of the non-blank values in `input`. The first element of the returned dataset is the mean and the second is the standard deviation.

This function will calculate both values in one pass over the dataset. Hence when both the mean and standard deviation of a dataset are necessary, this function is much more efficient than calling `gal_statistics_mean` and `gal_statistics_std` separately.

`gal_data_t *` [Function]
`gal_statistics_median` (*gal_data_t* `*input`, *int* `inplace`)

Return a single-element dataset containing the median of the non-blank values in `input`. The numerical datatype of the output is the same as `input`.

Calculating the median involves sorting the dataset and removing blank values, for better performance (and less memory usage), you can give a non-zero value to the

inplace argument. In this case, the sorting and removal of blank elements will be done directly on the input dataset. However, after this function the original dataset may have changed (if it wasn't sorted or had blank values).

size_t [Function]
gal_statistics_quantile_index (*size_t* size, *double* quantile)
Return the index of the element that has a quantile of quantile assuming the dataset has size elements.

size_t [Function]
gal_statistics_quantile (*gal_data_t* *input, *double* quantile, *int* inplace)
Return a single-element dataset containing the value with in a quantile quantile of the non-blank values in input. The numerical datatype of the output is the same as input. See gal_statistics_median for a description of inplace.

size_t [Function]
gal_statistics_quantile_function_index (*gal_data_t* *input, *gal_data_t* *value, *int* inplace)
Return the index of the quantile function (inverse quantile) of input at value. In other words, this function will return the index of the nearest element (of a sorted and non-blank) input to value. See gal_statistics_median for a description of inplace.

gal_data_t * [Function]
gal_statistics_quantile_function (*gal_data_t* *input, *gal_data_t* *value, *int* inplace)
Return a single-element (double or float64) dataset containing the quantile function of the non-blank values in input at value. In other words, this function will return the quantile of value in input. See gal_statistics_median for a description of inplace.

gal_data_t * [Function]
gal_statistics_mode (*gal_data_t* *input, *float* mirrordist, *int* inplace)
Return a four-element (double or float64) dataset that contains the mode of the input distribution. This function implements the non-parametric algorithm to find the mode that is described in Appendix C of Akhlaghi and Ichikawa [2015] (https://arxiv.org/abs/1505.01664).

In short it compares the actual distribution and its "mirror distribution" to find the mode. In order to be efficient, you can determine how far the comparison goes away from the mirror through the mirrordist parameter (think of it as a multiple of sigma/error). See gal_statistics_median for a description of inplace.

The output array has the following elements (in the given order, note that counting in C starts from 0).

```
array[0]: mode
array[1]: mode quantile.
array[2]: symmetricity.
array[3]: value at the end of symmetricity.
```

`gal_data_t *` [Function]
`gal_statistics_mode_mirror_plots` (*gal_data_t* `*input`, *gal_data_t* `*value`, *size_t* `numbins`, *int* `inplace`, *double* `*mirror_val`)

Make a mirrored histogram and cumulative frequency plot (with `numbins`) with the mirror distribution of the `input` with a value at `value`.

The output is a list of data structures (see Section 10.3.7.9 [List of `gal_data_t`], page 267): the first is the bins with one bin at the mirror point, the second is the histogram with a maximum of one and the third is the cumulative frequency plot (with a maximum of one).

`int` [Function]
`gal_statistics_is_sorted` (*gal_data_t* `*input`)

Return the respective sort macro (see above) for the `input` dataset.

`void` [Function]
`gal_statistics_sort_increasing` (*gal_data_t* `*input`)

Sort the input dataset (in place) in an increasing order.

`void` [Function]
`gal_statistics_sort_decreasing` (*gal_data_t* `*input`)

Sort the input dataset (in place) in a decreasing order.

`gal_data_t *` [Function]
`gal_statistics_no_blank_sorted` (*gal_data_t* `*input`, *int* `inplace`)

Remove all the blanks and sort the input dataset. If `inplace` is non-zero this will happen on the input dataset (and the output will be the same as the input). However, if `inplace` is zero, this function will allocate a new copy of the dataset that is sorted and has no blank values.

`gal_data_t *` [Function]
`gal_statistics_regular_bins` (*gal_data_t* `*input`, *gal_data_t* `*inrange`, *size_t* `numbins`, *double* `onebinstart`)

Generate an array of regularly spaced elements as a 1D array (column) of type `double` (i.e., `float64`, it has to be double to account for small differences on the bin edges). The input arguments are discribed below

`input` The dataset you want to apply the bins to. This is only necessary if the range argument is not complete, see below. If `inrange` has all the necessary information, you can pass a `NULL` pointer for this.

`inrange` This dataset keeps the desired range along each dimension of the input data structure, it has to be in `float` (i.e., `float32`) type.

- If you want the full range of the dataset (in any dimensions, then just set `inrange` to `NULL` and the range will be specified from the minimum and maximum value of the dataset (`input` cannot be `NULL` in this case).
- If there is one element for each dimension in range, then it is viewed as a quantile (Q), and the range will be: 'Q to 1-Q'.

- If there are two elements for each dimension in range, then they are assumed to be your desired minimum and maximum values. When either of the two are NaN, the minimum and maximum will be calculated for it.

`numbins` The number of bins: must be larger than 0.

`onebinstart`
A desired value for onebinstart. Note that with this option, the bins won't start and end exactly on the given range values, it will be slightly shifted to accommodate this request.

`gal_data_t *` [Function]
`gal_statistics_histogram` (*gal_data_t* `*input`, *gal_data_t* `*bins`, *int* `normalize`, *int* `maxone`)

Make a histogram of all the elements in the given dataset with bin values that are defined in the `inbins` structure (see `gal_statistics_regular_bins`). `inbins` is not mandatory, if you pass a `NULL` pointer, the bins structure will be built within this function based on the `numbins` input. As a result, when you have already defined the bins, `numbins` is not used.

`gal_data_t *` [Function]
`gal_statistics_cfp` (*gal_data_t* `*input`, *gal_data_t* `*bins`, *int* `normalize`)

Make a cumulative frequency plot (CFP) of all the elements in `input` with bin values that are defined in the `bins` structure (see `gal_statistics_regular_bins`).

The CFP is built from the histogram: in each bin, the value is the sum of all previous bins in the histogram. Thus, if you have already calculated the histogram before calling this function, you can pass it onto this function as the data structure in `bins->next` (see `List of gal_data_t`). If `bin->next!=NULL`, then it is assumed to be the histogram. If it is `NULL`, then the histogram will be calculated internally and freed after the job is finished.

When a histogram is given and it is normalized, the CFP will also be normalized (even if the normalized flag is not set here): note that a normalized CFP's maximum value is 1.

`gal_data_t *` [Function]
`gal_statistics_sigma_clip` (*gal_data_t* `*input`, *float* `multip`, *float* `param`, *int* `inplace`, *int* `quiet`)

Apply σ-clipping on a given dataset and return a dataset that contains the results. For a description of σ-clipping see Section 7.1.2 [Sigma clipping], page 149. `multip` is the multiple of the standard deviation (σ that is used to define outliers in each round of clipping.

The role of `param` is determined based on its value. If `param` is larger than `1` (one), it must be an integer and will be interpretted as the number clips to do. If it is less than `1` (one), it is interpretted as the tollerance level to stop the iteration.

The output dataset has the following elements:

```
array[0]: Number of points used.
array[1]: Median.
```

```
array[2]: Mean.
array[3]: Standard deviation.
```

10.3.19 Binary datasets (binary.h)

Binary datasets only have two (usable) values: 0 (also known as background) or 1 (also known as foreground). They are created after some binary classificataion is applied to the dataset. The most common is thresholding: for example in an image, pixels with a value above the threshold are given a value of 1 and those with a value less than the threshold are assigned a value of 0.

Since there is only two values, in the processing of binary images, you are usually concered with the positioning of an element and its vicinity (neighbors). When a dataset has more than one dimension, multiple classes of immediate neighbors (that are touching the element) can be defined for each data-element. To separate these different classes of immediate neighbors, we define *connectivity*.

The classification is done by the distance from element center to the neighbor's center. The nearest immediate neighbors have a connectivity of 1, the second nearest class of neighbors have a connectivity of 2 and so on. In total, the largest possible connectivity for data with `ndim` dimensions is `ndim`. For example in a 2D dataset, 4-connected neighbors (that share an edge and have a distance of 1 pixel) have a connectivity of 1. The other 4 neighbors that only share a vertice (with a distance of $\sqrt{2}$ pixels) have a connectivity of 2. Conventionally, the class of connectivity-2 neighbors also includes the connectivity 1 neighbors, so for example we call them 8-connected neighbors in 2D datasets.

Ideally, one bit is sufficient for each element of a binary dataset. However, CPUs are not designed to work on individual bits, the smallest unit of memory addresses is a byte (containing 8 bits on modern CPUs). Therefore, in Gnuastro, the type used for binary dataset is `uint8_t` (see Section 4.4 [Numeric data types], page 64). Although it does take 8-times more memory, this choice offers much better performance and the some extra (useful) features.

The advantage of using a full byte for each element of a binary dataset is that you can also have other values (that will be ignored in the processing). One such common "other" value in real datasets is a blank value (to mark regions that should not be processed because there is no data). The constant `GAL_BLANK_UINT8` value must be used in these cases (see Section 10.3.4 [Library blank values (`blank.h`)], page 243). Another is some temporary value(s) that can be given to a processed pixel to avoid having another copy of the dataset as in `GAL_BINARY_TMP_VALUE` that is described below.

`GAL_BINARY_TMP_VALUE` [Macro]

The functions described below work on a `uint8_t` type dataset with values of 1 or 0 (no other pixel will be touched). However, in some cases, it is necessary to put temporary values in each element during the processing of the functions. This temporary value has a special meaning for the operation and will be operated on. So if your input datasets have values other than 0 and 1 that you don't want these functions to work on, be sure they are not equal to this macro's value. Note that this value is also different from `GAL_BLANK_UINT8`, so your input datasets may also contain blank elements.

gal_data_t * [Function]
gal_binary_erode (*gal_data_t* *input, *size_t* num, *int* connectivity, *int* inplace)

Do num erosions on the connectivity-connected neighbors of input (see above for the defintion of connectivity).

If inplace is non-zero *and* the input's type is GAL_TYPE_UINT8, then the erosion will be done within the input dataset and the returned pointer will be input. Otherwise, input is copied (and converted if necessary) to GAL_TYPE_UINT8 and erosion will be done on this new dataset which will also be returned. This function will only work on the elements with a value of 1 or 0. It will leave all the rest unchanged.

Erosion (inverse of dilation) is an operation in mathematical morphology where each foreground pixel that is touching a background pixel is flipped (changed to background). The connectivity value determines the definition of "touching". Erosion will thus decrease the area of the foregound regions by one layer of pixels.

gal_data_t * [Function]
gal_binary_dilate (*gal_data_t* *input, *size_t* num, *int* connectivity, *int* inplace)

Do num dilations on the connectivity-connected neighbors of input (see above for the defintion of connectivity). For more on inplace and the output, see gal_binary_erode.

Dilation (inverse of erosion) is an operation in mathematical morphology where each background pixel that is touching a foreground pixel is flipped (changed to foreground). The connectivity value determines the definition of "touching". Dilation will thus increase the area of the foregound regions by one layer of pixels.

gal_data_t * [Function]
gal_binary_open (*gal_data_t* *input, *size_t* num, *int* connectivity, *int* inplace)

Do num openings on the connectivity-connected neighbors of input (see above for the defintion of connectivity). For more on inplace and the output, see gal_binary_erode.

Opening is an operation in mathematical morphology which is defined as erosion followed by dilation (see above for the definitions of erosion and dilation). Opening will thus remove the outer structure of the foreground. In this implementation, num erosions are going to be applied on the dataset, then num dilations.

size_t [Function]
gal_binary_connected_components (*gal_data_t* *binary, *gal_data_t* **out, *int* connectivity)

Return the number of connected components in binary through the breadth first search algorithm (finding all pixels belonging to one component before going on to the next). Connection between two pixels is defined based on the value to connectivity. out is a dataset with the same size as binary with GAL_TYPE_INT32 type. Every pixel in out will have the label of the connected component it belongs to. The labeling of connected components starts from 1, so a label of zero is given to the input's background pixels.

When `*out!=NULL` (its space is already allocated), it will be cleared (to zero) at the start of this function. Otherwise, when `*out==NULL`, the necessary dataset to keep the output will be allocated by this function.

`binary` must have a type of `GAL_TYPE_UINT8`, otherwise this function will abort with an error. Other than blank pixels (with a value of `GAL_BLANK_UINT8` defined in Section 10.3.4 [Library blank values (`blank.h`)], page 243), all other non-zero pixels in `binary` will be considered as foreground (and will be labeled). Blank pixels in the input will also be blank in the output.

`gal_data_t *` [Function]
`gal_binary_connected_adjacency_matrix` (*gal_data_t* `*adjacency`, *size_t* `*numconnected`)

Find the number of connected labels and new labels based on an adjacency matrix, which must be a square binary array (type `GAL_TYPE_UINT8`). The returned dataset is a list of new labels for each old label. In other words, this function will find the objects that are connected (possibly through a third object) and in the output array, the respective elements for all input labels is going to have the same value. The total number of connected labels is put into the space that `numconnected` points to.

An adjacency matrix defines connection between two labels. For example, let's assume we have 5 labels and we know that labels 1 and 5 are connected to label 3, but are not connected with each other. Also, labels 2 and 4 are not touching any other label. So in total we have 3 final labels: one combined object (merged from labels 1, 3, and 5) and the initial labels 2 and 4. The input adjacency matrix would look like this (note the extra row and column for a label 0 which is ignored):

```
                 INPUT                             OUTPUT
                 =====                             ======
               in_lab  1  2  3  4  5   |
                                       |       numconnected = 3
                    0  0  0  0  0  0   |
    in_lab 1 -->    0  0  0  1  0  0   |
    in_lab 2 -->    0  0  0  0  0  0   |  Returned: new labels for the
    in_lab 3 -->    0  1  0  0  0  1   |       5 initial objects
    in_lab 4 -->    0  0  0  0  0  0   |   | 0 | 1 | 2 | 1 | 3 | 1 |
    in_lab 5 -->    0  0  0  1  0  0   |
```

Although the adjacency matrix as used here is symmetric, currently this function assumes that it is filled on both sides of the diagonal.

`void` [Function]
`gal_binary_fill_holes` (*gal_data_t* `*input`)

Fill all the holes that are bounded within a 4-connected region of the binary `input` dataset. This function currently only works on a 2D dataset.

10.3.20 Convolution functions (`convolve.h`)

Convolution is a very common operation during data analysis and is thoroughly described as part of Gnuastro's Section 6.3 [Convolve], page 117, program which is fully devoted to this job. Because of the complete introduction that was presented there, we will directly skip onto the currently available convolution functions in Gnuastro's library.

As of this version, only spatial domain convolution is available in Gnuastro's libraries. We haven't had the time to liberate the frequency domain function convolution and deconvolution functions that are available in the Convolve program[21].

`gal_data_t *` [Function]
`gal_convolve_spatial` (*gal_data_t* `*tiles`, *gal_data_t* `*kernel`, *size_t* `numthreads`, *int* `edgecorrection`, *int* `convoverch`)

Convolve each node of the list of `tiles` (see Section 10.3.7.9 [List of `gal_data_t`], page 267, and Section 10.3.13 [Tessellation library (`tile.h`)], page 288) with `kernel` using `numthreads`. When `edgecorrection` is non-zero, it will correct for the edge dimming effects as discussed in Section 6.3.1.2 [Edges in the spatial domain], page 119.

To create a tessellation that fully covers an input image, you may use `gal_tile_full`, or `gal_tile_full_two_layers` to also define channels over your input dataset. These functions are discussed in Section 10.3.13.2 [Tile grid], page 294. You may then pass the list of tiles to this function. This is the recommended way to call this function because spatial domain convolution is slow and breaking the job into many small tiles and working on simultaneously on several threads can greatly speed up the processing.

If the tiles are defined within a channel (a larger tile), by default convolution will be done within the channel, so pixels on the edge of a channel will not be affected by their neighbors that are in another channel. See Section 4.6 [Tessellation], page 72, for the necessity of channels in astronomical data analysis. This behavior may be disabled when `convoverch` is non-zero. In this case, it will ignore channel borders (if they exist) and mix all pixels that cover the kernel within the dataset.

`void` [Function]
`gal_convolve_spatial_correct_ch_edge` (*gal_data_t* `*tiles`, *gal_data_t* `*kernel`, *size_t* `numthreads`, *int* `edgecorrection`, *gal_data_t* `*tocorrect`)

Correct the edges of channels in an already convolved image when it was initially convolved with `gal_convolve_spatial` and `convoverch==0`. In that case, strong boundaries might exist on the channel edges. So if you later need to remove those boundaries at later steps of your processing, you can call this function. It will only do convolution on the tiles that are near the edge and were effected by the channel borders. Other pixels in the image will not be touched. Hence, it is much faster.

10.3.21 Interpolation (`interpolate.h`)

During data anlysis, it often happens that parts of the data cannot be given a value. For example your image was saturated due to a very bright start and you have to mask that star's footprint. One other common situation in Gnuastro is when we do processing on tiles (for example to estimate the Sky value and its Standard deviation, see Section 7.1.3 [Sky value], page 150). Some tiles must not be used for the estimation of the Sky value, for example because they cover a large galaxy. So we need to fill them in with blank values. But ultimately, we need a Sky value for every pixel. This job (assigning a value to blank element(s) based on their nearby neighbors with a value) is known as interpolation.

[21] Hence any help would be greatly appreciated.

There are many ways to do interpolation, but (mainly due to lack of time), currently Gnuastro only contains the (median of) nearest-neighbor method. The power of this method of interpolation is its non-parametric nature. The produced values are also always within the range of the known values and strong outliers do not get created. We will hopefully implement other methods too (wrappers around the GNU Scientific Library's very complete set of functions), but currently the developers are too busy. So if you do have the chance to help your contribution would be very welcome and appreciated.

`gal_data_t *` [Function]
`gal_interpolate_close_neighbors` (*gal_data_t* `*input`, *struct gal_tile_two_layer_params* `*tl`, *size_t* `numneighbors`, *size_t* `numthreads`, *int* `onlyblank`, *int* `aslinkedlist`)

Interpolate the values in the image using the median value of their `numneighbors` closest neighbors. If `onlyblank` is non-zero, then only blank elements will be interpolated and pixels that already have a value will be left untouched. This function is multi-threaded and will run on `numthreads` threads (see `gal_threads_number` in Section 10.3.2 [Multithreaded programming (`threads.h`)], page 235).

`tl` is Gnuastro's two later tessellation structure used to define tiles over an image and is fully described in Section 10.3.13.2 [Tile grid], page 294. When `tl!=NULL`, then it is assumed that the `input->array` contains one value per tile and interpolation will respect certain tessellation properties, for example to not interpolate over channel borders.

If several datasets have the same set of blank values, you don't need to call this function multiple times. When `aslinkedlist` is non-zero, then `input` will be seen as a Section 10.3.7.9 [List of `gal_data_t`], page 267, and for all the same neighbor position checking will be done for all the datasets in the list. Ofcourse, the values for each dataset will be different, so a different value will be written in the each dataset, but the neighbor checking that is the most CPU intensive part will only be done once.

10.3.22 Git wrappers (`git.h`)

Git is one of the most common tools for version control and it can often be useful during development, for example see `COMMIT` keyword in Section 4.9 [Output headers], page 78. At installation time, Gnuastro will also check for the existence of libgit2, and store the value in the `GAL_CONFIG_HAVE_LIBGIT2`, see Section 10.3.1 [Configuration information (`config.h`)], page 234, and Section 3.1.2 [Optional dependencies], page 28. `gnuastro/git.h` includes `gnuastro/config.h` internally, so you won't have to include both for this macro.

`char *` [Function]
`gal_git_describe` ( )

When libgit2 is present and the program is called within a directory that is version controlled, this function will return a string containing the commit description (similar to Gnuastro's unofficial version number, see Section 1.5 [Version numbering], page 5). If there are uncommitted changes in the running directory, it will add a '`-dirty`' prefix to the description. When there is no tagged point in the previous commit, this function will return a uniquely abbreviated commit object as fallback. This function is used for generating the value of the `COMMIT` keyword in Section 4.9 [Output headers], page 78. The output string is similar to the output of the following command:

```
$ git describe --dirty --always
```

Space for the output string is allocated within this function, so after using the value you have to `free` the output string. If libgit2 is not installed or the program calling this function is not within a version controlled directory, then the output will be the `NULL` pointer.

10.4 Library demo programs

In this final section of Chapter 10 [Library], page 223, we give some example Gnuastro programs to demonstrate various features in the library. All these programs have been tested and once Gnuastro is installed you can compile and run them with with Gnuastro's Section 10.2 [BuildProgram], page 230, program that will take care of linking issues. If you don't have any FITS file to experiment on, you can use those that are generated by Gnuastro after `make check` in the `tests/` directory, see Section 1.1 [Quick start], page 1.

10.4.1 Library demo - reading a FITS image

The following simple program demonstrates how to read a FITS image into memory and use the `void *array` pointer in of Section 10.3.5.1 [Generic data container (`gal_data_t`)], page 245. For easy linking/compilation of this program along with a first run see Section 10.2 [BuildProgram], page 230. Before running, also change the `filename` and `hdu` variable values to specify an existing FITS file and/or extension/HDU.

This is just intended to demonstrate how to use the `array` pointer of `gal_data_t`. Hence it doesn't do important sanity checks, for example in real datasets you may also have blank pixels. In such cases, this program will return a NaN value (see Section 6.1.3 [Blank pixels], page 100). So for general statistical information of a dataset, it is much better to use Gnuastro's Section 7.1 [Statistics], page 148, program which can deal with blank pixels any many other issues in a generic dataset.

```
#include <stdio.h>
#include <stdlib.h>
#include <gnuastro/fits.h> /* includes gnuastro's data.h and type.h */
#include <gnuastro/statistics.h>

int
main(void)
{
  size_t i;
  float *farray;
  double sum=0.0f;
  gal_data_t *image;
  char *filename="img.fits", *hdu="1";


  /* Read 'img.fits' (HDU: 1) as a float32 array. */
  image=gal_fits_img_read_to_type(filename, hdu, GAL_TYPE_FLOAT32, -1);
```

```
  /* Use the allocated space as a single precision floating
   * point array (recall that ‘image->array’ has ‘void *’
   * type, so it is not directly usable. */
  farray=image->array;


  /* Calculate the sum of all the values. */
  for(i=0; i<image->size; ++i)
    sum += farray[i];


  /* Report the sum. */
  printf("Sum of values in %s (hdu %s) is: %f\n",
         filename, hdu, sum);


  /* Clean up and return. */
  gal_data_free(image);
  return EXIT_SUCCESS;
}
```

10.4.2 Library demo - inspecting neighbors

The following simple program shows how you can inspect the neighbors of a pixel using the `GAL_DIMENSION_NEIGHBOR_OP` function-like macro that was introduced in Section 10.3.6 [Dimensions (`dimension.h`)], page 253. For easy linking/compilation of this program along with a first run see Section 10.2 [BuildProgram], page 230. Before running, also change the `filename` and `hdu` variable values to specify an existing FITS file and/or extension/HDU.

```
#include <stdio.h>
#include <gnuastro/fits.h>
#include <gnuastro/dimension.h>

int
main(void)
{
  double sum;
  float *array;
  size_t i, num, *dinc;
  gal_data_t *input=gal_fits_img_read_to_type("input.fits", "1",
                                              GAL_TYPE_FLOAT32, -1);

  /* To avoid the ‘void *’ pointer and have ‘dinc’. */
  array=input->array;
  dinc=gal_dimension_increment(input->ndim, input->dsize);

  /* Go over all the pixels. */
  for(i=0;i<input->size;++i)
```

```
    {
      num=0;
      sum=0.0f;
      GAL_DIMENSION_NEIGHBOR_OP( i, input->ndim, input->dsize,
                                 input->ndim, dinc,
                                 {++num; sum+=array[nind];} );
      printf("%zu: num: %zu, sum: %f\n", i, num, sum);
    }

  /* Clean up and return. */
  gal_data_free(input);
  return EXIT_SUCCESS;
}
```

10.4.3 Library demo - multi-threaded operation

The following simple program shows how to use Gnuastro to simplify spinning off threads and distributing different jobs between the threads. The relevant thread-related functions are defined in Section 10.3.2.2 [Gnuastro's thread related functions], page 237. For easy linking/compilation of this program, along with a first run, see Gnuastro's Section 10.2 [BuildProgram], page 230. Before running, also change the `filename` and `hdu` variable values to specify an existing FITS file and/or extension/HDU.

This is a very simple program to open a FITS image, distribute its pixels between different threads and print the value of each pixel and the thread it was assigned to. The actual operation is very simple (and would not usually be done with threads in a real-life program). It is intentionally chosen to put more focus on the important steps in spinning of threads and how the worker function (which is called by each thread) can identify the job-IDs it should work on.

For example, instead of an array of pixels, you can define an array of tiles or any other context-specific strcutures as separate targets. The important thing is that each action should have its own unique ID (counting from zero, as is done in an array in C). You can then follow the process below and use each thread to work on all the targets that are assigned to it. Recall that spinning-off threads is its self an expensive process and we don't want to spin-off one thread for each target (see the description of `gal_threads_dist_in_threads` in Section 10.3.2.2 [Gnuastro's thread related functions], page 237.

There are many (more complicated, real-world) examples of using `gal_threads_spin_off` in Gnuastro's actual source code, you can see them by searching for the `gal_threads_spin_off` function from the top source (after unpacking the tarball) directory (for example with this command):

```
$ grep -r gal_threads_spin_off ./
```

The code of this demonstration program is shown below. This program was also built and run when you ran `make check` during the building of Gnuastro (`tests/lib/multithread.c`, so it is already tested for your system and you can safely use it as a guide.

```
#include <stdio.h>
#include <stdlib.h>
```

```
#include "gnuastro/fits.h"
#include "gnuastro/threads.h"


/* This structure can keep all information you want to pass onto the
 * worker function on each thread. */
struct params
{
  gal_data_t *image;            /* Dataset to print values of. */
};




/* This is the main worker function which will be called by the
 * different threads. ‘gal_threads_params’ is defined in
 * ‘gnuastro/threads.h’ and contains the pointer to the paramter we
 * want. Note that the input argument and returned value of this
 * function always must have ‘void *’ type. */
void *
worker_on_thread(void *in_prm)
{
  /* Low-level definitions to be done first. */
  struct gal_threads_params *tprm=(struct gal_threads_params *)in_prm;
  struct params *p=(struct params *)tprm->params;


  /* Subsequent definitions. */
  float *array=p->image->array;
  size_t i, index, *dsize=p->image->dsize;


  /* Go over all the actions(pixels in this case that were assigned
   * to this thread. */
  for(i=0; tprm->indexs[i] != GAL_BLANK_SIZE_T; ++i)
    {
      /* For easy reading. */
      index = tprm->indexs[i];


      /* Print the information. */
      printf("(%zu, %zu) on thread %zu: %g\n", index%dsize[1]+1,
             index/dsize[1]+1, tprm->id, array[index]);
    }


  /* Wait for all the other threads to finish, then return. */
```

```
  if(tprm->b) pthread_barrier_wait(tprm->b);
  return NULL;
}




/* High-level function (called by the operating system). */
int
main(void)
{
  struct params p;
  char *filename="input.fits", *hdu="1";
  size_t numthreads=gal_threads_number();


  /* Read the image into memory as a float32 data type. We are using
   * ‘-1’ for ‘minmapsize’ to ensure that the image is read into
   * memory. */
  p.image=gal_fits_img_read_to_type(filename, hdu, GAL_TYPE_FLOAT32,
                                    -1);


  /* Print some basic information before the actual contents: */
  printf("Pixel values of %s (HDU: %s) on %zu threads.\n", filename,
         hdu, numthreads);
  printf("Used to check the compiled library’s capability in opening "
         "a FITS file, and also spinning-off threads.\n");


  /* A small sanity check: this is only intended for 2D arrays (to
   * print the coordinates of each pixel). */
  if(p.image->ndim!=2)
    {
      fprintf(stderr, "only 2D images are supported.");
      exit(EXIT_FAILURE);
    }


  /* Spin-off the threads and do the processing on each thread. */
  gal_threads_spin_off(worker_on_thread, &p, p.image->size, numthreads);


  /* Clean up and return. */
  gal_data_free(p.image);
  return EXIT_SUCCESS;
}
```

10.4.4 Library demo - reading and writing table columns

Tables are some of the most common inputs to, and outputs of programs. This section contains a small program for reading and writing tables using the constructs described in Section 10.3.11 [Table input output (`table.h`)], page 281. For easy linking/compilation of this program, along with a first run, see Gnuastro's Section 10.2 [BuildProgram], page 230. Before running, also set the following file and column names in the first two lines of `main`. The input and output names may be `.txt` and `.fits` tables, `gal_table_read` and `gal_table_write` will be able to write to both formats. For plain text tables see see Section 4.5.2 [Gnuastro text table format], page 69.

This example program reads three columns from a table. The first two columns are selected by their name (`NAME1` and `NAME2`) and the third is selected by its number: column 10 (counting from 1). Gnuastro's column selection is discussed in Section 4.5.3 [Selecting table columns], page 71. The first and second columns can be any type, but this program will convert them to `int32_t` and `float` for its internal usage respectively. However, the third column must be double for this program. So if it isn't the program will abort with an error. Having the columns in memory, it will print them out along with their sum (just a simple application, you can do what ever you want at this stage). Reading the table finishes here.

The rest of the program is a demonstration of writing a table. While parsing the rows, this program will change the first column (to be counters) and multiply the second by 10 (to the output may be different). Then it will define the order of the output columns by setting the `next` element (to create a Section 10.3.7.9 [List of `gal_data_t`], page 267). Before writing, this function will also set names for the columns (units and comments can be defined in a similar manner). Writing the columns to a file is then done through a simple call to `gal_table_write`.

The operations that are shows in this example program are not necessary all the time. For example, in many cases, you know the numerical data type of the column before writing the program (see Section 4.4 [Numeric data types], page 64), so the type checkings and copyings won't be necessary.

```
#include <stdio.h>
#include <stdlib.h>

#include <gnuastro/table.h>

int
main(void)
{
  /* File names and column names (which may also be numbers). */
  char *c1_name="NAME1", *c2_name="NAME2", *c3_name="10";
  char *inname="input.fits", *hdu="1", *outname="out.fits";

  /* Internal parameters. */
  float *array2;
  double *array3;
  int32_t *array1;
  size_t i, counter=0;
```

```
gal_data_t *c1, *c2;
gal_data_t tmp, *col, *columns;
gal_list_str_t *column_ids=NULL;

/* Define the columns to read. */
gal_list_str_add(&column_ids, c1_name, 0);
gal_list_str_add(&column_ids, c2_name, 0);
gal_list_str_add(&column_ids, c3_name, 0);

/* The columns were added in reverse, so correct it. */
gal_list_str_reverse(&column_ids);

/* Read the desired columns. */
columns = gal_table_read(inname, hdu, column_ids,
                         GAL_TABLE_SEARCH_NAME, 1, -1);

/* Go over the columns, we'll assume that you don't know their type
 * a-priori, so we'll check  */
counter=1;
for(col=columns; col!=NULL; col=col->next)
  switch(counter++)
    {
    case 1:              /* First column: we want it as int32_t. */
      c1=gal_data_copy_to_new_type(col, GAL_TYPE_INT32);
      array1 = c1->array;
      break;

    case 2:              /* Second column: we want it as float.  */
      c2=gal_data_copy_to_new_type(col, GAL_TYPE_FLOAT32);
      array2 = c2->array;
      break;

    case 3:              /* Third column: it MUST be double.     */
      if(col->type!=GAL_TYPE_FLOAT64)
        {
          fprintf(stderr, "Column %s must be float64 type, it is "
                  "%s", c3_name, gal_type_name(col->type, 1));
          exit(EXIT_FAILURE);
        }
      array3 = col->array;
      break;
    }

/* As an example application we'll just print them out. In the
 * meantime (just for a simple demonstration), change the first
 * array value to the counter and multiply the second by 10. */
for(i=0;i<c1->size;++i)
```

```
    {
      printf("%zu: %d + %f + %f = %f\n", i+1, array1[i], array2[i],
             array3[i], array1[i]+array2[i]+array3[i]);
      array1[i]  = i+1;
      array2[i] *= 10;
    }

  /* Link the first two columns as a list. */
  c1->next = c2;
  c2->next = NULL;

  /* Set names for the columns and write them out. */
  c1->name = "COUNTER";
  c2->name = "VALUE";
  gal_table_write(c1, NULL, GAL_TABLE_FORMAT_BFITS, outname, 1);

  /* The names weren't allocated, so to avoid cleaning-up problems,
   * we'll set them to NULL. */
  c1->name = c2->name = NULL;

  /* Clean up and return.  */
  gal_data_free(c1);
  gal_data_free(c2);
  gal_list_data_free(columns);
  gal_list_str_free(column_ids, 0); /* strings weren't allocated. */
  return EXIT_SUCCESS;
}
```

11 Developing

The basic idea of GNU Astronomy Utilities is for an interested astronomer to be able to easily understand the code of any of the programs or libraries, be able to modify the code if s/he feels there is an improvement and finally, to be able to add new programs or libraries for their own benefit, and the larger community if they are willing to share it. In short, we hope that at least from the software point of view, the "obscurantist faith in the expert's special skill and in his personal knowledge and authority" can be broken, see Section 1.2 [Science and its tools], page 2. With this aim in mind, Gnuastro was designed to to have a very basic, simple, and easy to understand architecture for any interested inquirer.

This chapter starts with very general design choices, in particular Section 11.1 [Why C programming language?], page 320, and Section 11.2 [Program design philosophy], page 322. It will then get a little more technical about the Gnuastro code and file/directory structure in Section 11.3 [Coding conventions], page 323, and Section 11.4 [Program source], page 326. Section 11.4.2 [The TEMPLATE program], page 329, discusses a minimal (and working) template to help in creating new programs or easier learning of a program's internal structure. Some other general issues about documentation, building and debugging are then discussed. This chapter concludes with how you can learn about the development and get involved in Section 11.9 [Gnuastro project webpage], page 333, Section 11.10 [Developing mailing lists], page 335, and Section 11.11 [Contributing to Gnuastro], page 335.

11.1 Why C programming language?

Currently the programming language that is most commonly used in scientific applications is C++[1], Python[2], and Julia[3] (which is a newcomer but swiftly gaining ground). One of the main reasons behind this choice is their high-level abstractions. However, GNU Astronomy Utilities is fully written in the C programming language[4]. The reasons can be summarized with simplicity, portability and efficiency/speed. All three are very important in a scientific software and we will discuss them below.

Simplicity can best be demonstrated in a comparison of the main books of C++ and C. The "C programming language"[5] book, written by the authors of C, is only 286 pages and covers a very good fraction of the language, it has also remained unchanged from 1988. C is the main programming language of nearly all operating systems and there is no plan of any significant update. On the other hand, the most recent "C++ programming language"[6] book, also written by its author, has 1366 pages and its fourth edition came out in 2013! As discussed in Section 1.2 [Science and its tools], page 2, it is very important for other scientists to be able to readily read the code of a program at their will with minimum requirements.

In C++, inheriting objects in the object oriented programming paradigm and their internal functions make the code very easy to write for a programmer who is deeply invested in

---

[1] https://isocpp.org/

[2] https://www.python.org/

[3] https://julialang.org/

[4] https://en.wikipedia.org/wiki/C_(programming_language)

[5] Brian Kernighan, Dennis Ritchie. *The C programming language*. Prentice Hall, Inc., Second edition, 1988. It is also commonly known as K&R and is based on the ANSI C and ISO C90 standards.

[6] Bjarne Stroustrup. *The C++ programming language*. Addison-Wesley Professional; 4 edition, 2013.

those objects and understands all their relations well. But it simultaneously makes reading the program for a first time reader (a curious scientist who wants to know only how a small step was done) extremely hard. Before understanding the methods, the scientist has to invest a lot of time and energy in understanding those objects and their relations. But in C, everything is done with basic language types for example `ints` or `floats` and their pointers to define arrays. So when an outside reader is only interested in one part of the program, that part is all they have to understand.

Recently it is also becoming common to write scientific software in Python, or a combination of it with C or C++. Python is a high level scripting language which doesn't need compilation. It is very useful when you want to do something on the go and don't want to be halted by the troubles of compiling, linking, memory checking, etc. When the datasets are small and the job is temporary, this ability of Python is great and is highly encouraged. A very good example might be plotting, in which Python is undoubtedly one of the best.

But as the data sets increase in size and the processing becomes more complicated, the speed of Python scripts significantly decrease. So when the program doesn't change too often and is widely used in a large community, mostly on large data sets (like astronomical images), using Python will waste a lot of valuable research-hours. It is possible to wrap C or C++ functions with Python to fix the speed issue. But this adds to complexity, because the interested scientist will now have to master two programming languages and their connection (which is not trival).

Like C++, Python is object oriented, so as explained above, it needs a high level of experience with that particular program to fully understand its inner workings. To make things worse, since it is mainly for fast and on the go programming[7], it can undergo significant changes. One recent example is how Python 2.x and Python 3.x are not compatible. Lots of research teams that invested heavily in Python 2.x cannot benefit from Python 3.x or future versions any more. Some converters are available, but since they are automatic, lots of complications might arise in the conversion. Thus, re-writing all the changes would be the only truly reliable option. If a research project begins using Python 3.x today, there is no telling how compatible their investments will be when Python 4.x or 5.x will come out. This stems from the core principles of Python, which are very useful when you look in the 'on the go' basis as described before and not future usage. Future-proof code (as long as current operating systems will be used) is written in C.

The portability of C is best demonstrated by the fact that both C++ and Python are part of the C-family of programming languages which also include Julia, Java, Perl, and many other languages. C libraries can be immediately included in C++, and it is easy to write wrappers for them in all C-family programming languages. This will allow other scientists to benefit from C libraries using any C-family language that they prefer. As a result, Gnuastro's library is already usable in C and C++, and wrappers will be[8] added for higher-level languages like Python, Julia and Java.

The final reason was speed. This is another very important aspect of C which is not independent of simplicity (first reason discussed above). The abstractions provided by the higher-level languages (which also makes learning them harder for a newcomer) comes at

[7] Note that Python is good for fast programming, not fast programs.

[8] `http://savannah.gnu.org/task/?13786`

the cost of speed. Since C is a low-level language[9] (closer to the hardware), it is much less complex for both the human reader *and* the computer. The benefits of simplicity for a human were discussed above. Simplicity for the computer translates into more efficiently (faster) programs. This creates a much closer relation between the scientist/programmer (or their program) and the actual data and processing. The GNU coding standards[10] also encourage the use of C over all other languages when generality of usage and "high speed" is desired.

11.2 Program design philosophy

The core processing functions of each program (and all libraries) are written mostly with the basic ISO C90 standard. We do make lots of use of the GNU additions to the C language in the GNU C library, but these additional functions are mainly used in the user interface functions (reading your inputs and preparing them prior to or after the analysis). The actual algorithms, which most scientists would be more interested in, are much more closer to ISO C90. For this reason, program source files that deal with user interface issues and those doing the actual processing are clearly separated, see Section 11.4 [Program source], page 326. If anything particular to the GNU C library is used in the processing functions, it is explained in the comments in between the code.

Similar to GNU Coreutils, all the Gnuastro programs provide very low level operations. This enables you to use shell scripting languages (for example GNU Bash) to operate on a large number of files or do very complex things through the creative combinations of these tools that the authors had never dreamed of. We have put a few simple examples in Chapter 2 [Tutorials], page 14.

For example all the analysis output can be saved as ASCII tables which can be fed into your favorite plotting program to inspect visually. Python's Matplotlib is very useful for fast plotting of the tables to immediately check your results. If you want to include the plots in a document, you can use the PGFplots package within LaTeX, no attempt is made to include such operations in Gnuastro. In short, Bash can act as a glue to connect the inputs and outputs of all these various Gnuastro programs (and other programs) in any fashion you please. Ofcourse, Gnuastro's programs are just front-ends to the main workhorse (Section 10.3 [Gnuastro library], page 233), allowing a user to create their own programs (for example with Section 10.2 [BuildProgram], page 230). So once the functions within programs become mature enough, they will be moved within the libraries for more general applications.

The advantage of this architecture is that the programs become small and transparent: the starting and finishing point of every program is clearly demarcated. For nearly all operations on a modern computer, the read/write speed is very insignificant compared to the actual processing a program does. Therefore the complexity which arises from sharing memory in a large application is simply not worth the speed gain. Gnuastro's design is heavily influenced from Eric Raymond's "The Art of Unix Programming"[11] which beauti-

---

[9] Low-level languages are those that directly operate the hardware like assembly languages. So C is actually a high-level language, but it can be considered one of the lowest-level languages among all high-level languages.

[10] `http://www.gnu.org/prep/standards/`

[11] Eric S. Raymond, 2004, *The Art of Unix Programming*, Addison-Wesley Professional Computing Series.

fully describes the design philosophy and practice which lead to the success of Unix-based operating systems[12].

11.3 Coding conventions

In Gnuastro, we try our best to follow the GNU coding standards. Added to those, Gnuastro defines the following conventions. It is very important for readability that the whole package follows the same convention.

- The code must be easy to read by eye. So when the order of several lines within a function does not matter (for example when defining variables at the start of a function). You should put the lines in the order of increasing length and group the variables with similar types such that this half-pyramid of declarations becomes most visible. If the reader is interested, a simple search will show them the variable they are interested in. However, this visual aid greatly helps in general inspections of the code and help the reader get a grip of the function's processing.
- When ever you see that the function cannot be fully displayed (vertically) in your monitor, this is a sign that it has probably become too long and should be broken up into multiple functions. 40 lines is usually a good reference. When the start and end of a function are clearly visible in one glance, the function is much more easier to understand. This is most important for low-level functions (which usually define a lot of variables). Low-level functions do most of the processing, they will also be the most interesting part of a program for an inquiring astronomer. This convention is less important for higher level functions that don't define too many variables and whose only purpose is to run the lower-level functions in a specific order and with checks.

  In general you can be very liberal in breaking up the functions into smaller parts, the GNU Compiler Collection (GCC) will automatically compile the functions as inline functions when the optimizations are turned on. So you don't have to worry about decreasing the speed. By default Gnuastro will compile with the `-O3` optimization flag.
- All Gnuastro hand-written text files (C source code, Texinfo documentation source, and version control commit messages) should not exceed **75** characters per line. Monitors today are certainly much wider, but with this limit, reading the functions becomes much more easier. Also for the developers, it allows multiple files (or multiple views of one file) to be displayed beside each other on wide monitors.

  Emacs's buffers are excellent for this capability, setting a buffer width of 80 with '`C-u 80 C-x 3`' will allow you to view and work on several files or different parts of one file using the wide monitors common today. Emacs buffers can also be used as a shell prompt and compile the program (with `M-x compile`), and 80 characters is the default width in most terminal emulators. If you use Emacs, Gnuastro sets the 75 character `fill-column` variable automatically for you, see cartouche below.

  For long comments you can use press `Alt-q` in Emacs to separate them into separate lines automatically. For long literal strings, you can use the fact that in C, two strings immediately after each other are concatenated, for example `"The first part, " "and the second part."`. Note the space character in the end of the first part. Since they are now separated, you can easily break a long literal string into several lines and adhere to the maximum 75 character line length policy.

[12] KISS principle: Keep It Simple, Stupid!

- The headers required by each source file (ending with `.c`) should be defined inside of it. All the headers a complete program needs should *not* be stacked in another header to include in all source files (for example `main.h`). Although most 'professional' programmers choose this single header method, Gnuastro is primarily written for professional/inquisitive astronomers (who are generally amateur programmers). The list of header files included provides valuable general information and helps the reader. `main.h` may only include the header file(s) that define types that the main program structure needs, see `main.h` in Section 11.4 [Program source], page 326. Those particular header files that are included in `main.h` can of course be ignored (not included) in separate source files.
- The headers should be classified (by an empty line) into separate groups:
  1. `#include <config.h>`: This must be the first code line (not commented or blank) in each source file *within Gnuastro*. It sets macros that the GNU Portability Library (Gnulib) will use for a unified environment (GNU C Library), even when the user is building on a system that doesn't use the GNU C library.
  2. The C library header files, for example `stdio.h`, `stdlib.h`, or `math.h`.
  3. Installed library header files, including Gnuastro's installed headers (for example `cfitsio.h` or `gsl/gsl_rng.h`, or `gnuastro/fits.h`).
  4. Gnuastro's internal headers (that are not installed), for example `gnuastro-internal/options.h`.
  5. For programs, the `main.h` file (which is needed by the next group of headers).
  6. That particular program's header files, for example `mkprof.h`, or `noisechisel.h`.

  As much as order does not matter when you include the header of each group, sort them by length, as described above.
- All function names, variables, etc should be in lower case. Macros, constant global `enum`s should be in upper case.
- Naming of exported header files, functions, variables, macros, and library functions, we adopt similar conventions to those used by the GNU Scientific Library (GSL)[13]. In particular, in order to avoid clashes with the names of functions and variables coming from other libraries the name-space '`gal_`' is prefixed to them. GAL stands for *G*NU *A*stronomy *L*ibrary.
- All installed header files should be in the `lib/gnuastro` directory (under the top Gnuastro source directory). After installation, they will be put in the `$prefix/include/gnuastro` directory (see Section 3.3.1.2 [Installation directory], page 39, for `$prefix`). Therefore with this convention Gnuastro's headers can be included in internal (to Gnuastro) and external (a library user) source files with the same line

  ```
  # include <gnuastro/headername.h>
  ```

  Note that the GSL convention for header file names is `gsl_specialname.h`, so your include directive must be something like `#include <gsl/gsl_specialname.h>`, the header file names are not prefixed with a '`gal_`'. However, Gnuastro doesn't follow this guideline because of the repeated `gsl` in the include directive (which canbe confusing

[13] https://www.gnu.org/software/gsl/design/gsl-design.html#SEC15

and cause bugs). All Gnuastro (and GSL) headers must be located within a unique directory and will not be mixed with other headers. Therefore the '`gsl_`' prefix to the header file names is redundant[14].

- All installed functions and variables should also include the base-name of the file in which they are defined as prefix, using underscores to separate words[15]. The same applies to exported macros, but in upper case. For example in Gnuastro's top source directory, the prototype of function `gal_box_border_from_center` is in `lib/gnuastro/box.h`, and the macro `GAL_POLYGON_MAX_CORNERS` is defined in `lib/gnuastro/polygon.h`.

  This is necessary to give any user (who is not familiar with the library structure) the ability to follow the code of a function that is not in the program or library they are reading. This convention does make the function names longer (a little harder to write), but the extra documentation it provides convention plays an important role in Gnuastro and is worth the cost.
- There should be no trailing white space in a line. To do this automatically every time you save a file in Emacs, add the following line to your `~/.emacs` file.

  ```
  (add-hook 'before-save-hook 'delete-trailing-whitespace)
  ```
- There should be no tabs in the indentation[16].
- Individual, contextually similar, functions in a source file are separated by 5 blank lines to be easily seen to be related in a group when parsing the source code by eye. In Emacs you can use `CTRL-u 5 CTRL-o`.
- One group of contextually similar functions in a source file is separated from another with 20 blank lines. In Emacs you can use `CTRL-u 20 CTRL-o`. Each group of functions has short descriptive title of the functions in that group. This title is surrounded by asterisks (`*`) to make it clearly distinguishable. Such contextual grouping and clear title are very important for easily understanding the code.
- Always read the comments before the patch of code under it. Similarly, try to add as many comments as you can regarding every patch of code. Effectively, we want someone to get a good feeling of the steps, without having to read the C code and only by reading the comments. This follows similar principles as Literate programming (`https://en.wikipedia.org/wiki/Literate_programming`).

The last two conventions are not common and might benefit from a short discussion here. With a good experience in advanced text editor operations, the last two are redundant for a professional developer. However, recall that Gnuastro aspires to be friendly to un-familiar, and un-experienced (in programming) eyes. In other words, as discussed in Section 1.2 [Science and its tools], page 2, we want the code to appear welcoming to someone who is completely new to coding (and text editors) and only has a scientific curiosity.

---

[14] For GSL, this prefix has an internal technical application: GSL's architecture mixes installed and not-installed headers in the same directory. This prefix is used to identify their installation status. Therefore this filename prefix in GSL a technical internal issue (for developers, not users).

[15] The convention to use underscores to separate words, called "snake case" (or "snake_case"). This is also recommended by the GNU coding standards.

[16] If you use Emacs, Gnuastro's `.dir-locals.el` file will automatically never use tabs for indentation. To make this a default in all your Emacs sessions, you can add the following line to your `~/.emacs` file: `(setq-default indent-tabs-mode nil)`

Newcomers to coding and development, who are curious enough to venture into the code, will probably not be using (or have any knowledge of) advanced text editors. They will see the raw code in the webpage or on a simple text editor (like Gedit) as plain text. Trying to learn and understand a file with dense functions that are all spaced with one or two blank lines can be very taunting for a newcomer. But when they scroll through the file and see clear titles and meaningful spaces for similar functions, we are helping them find and focus on the part they are most interested in sooner and easier.

GNU Emacs, the recommended text editor: GNU Emacs is an extensible and easily customizable text editor which many programmers rely on for developing due to its countless features. Among them, it allows specification of certain settings that are applied to a single file or to all files in a directory and its sub-directories. In order to harmonize code coming from different contributors, Gnuastro comes with a `.dir-locals.el` file which automatically configures Emacs to satisfy most of the coding conventions above when you are using it within Gnuastro's directories. Thus, Emacs users can readily start hacking into Gnuastro. If you are new to developing, we strongly recommend this editor. Emacs was the first project released by GNU and is still one of its flagship projects. Some resources can be found at:

Official manual

At `https://www.gnu.org/software/emacs/manual/emacs.html`. This is a great and very complete manual which is being improved for over 30 years and is the best starting point to learn it. It just requires a little patience and practice, but rest assured that you will be rewarded. If you install Emacs, you also have access to this manual on the command-line with the following command (see Section 4.7.4 [Info], page 76).

```
$ info emacs
```

A guided tour of emacs

At `https://www.gnu.org/software/emacs/tour/`. A short visual tour of Emacs, officially maintained by the Emacs developers.

Unofficial mini-manual

At `https://tuhdo.github.io/emacs-tutor.html`. A shorter manual which contains nice animated images of using Emacs.

11.4 Program source

Besides the fact that all the programs share some functions that were explained in Chapter 10 [Library], page 223, everything else about each program is completely independent. Recall that Gnuastro is written for an active astronomer/scientist (not a passive one who just uses a software). It must thus be easily navigable. Hence there are fixed source files (that contain fixed operations) that must be present in all programs, these are discussed fully in Section 11.4.1 [Mandatory source code files], page 327. To easily understand the explanations in this section you can use Section 11.4.2 [The TEMPLATE program], page 329, which contains the bare minimum code for one working program. This

template can also be used to easily add new utilities: just copy and paste the directory and change `TEMPLATE` with your program's name.

11.4.1 Mandatory source code files

Some programs might need lots of source files and if there is no fixed convention, navigating them can become very hard for a new inquirer into the code. The following source files exist in every program's source directory (which is located in `bin/progname`). For small programs, these files are enough. Larger programs will need more files and developers are encouraged to define any number of new files. It is just important that the following list of files exist and do what is described here. When creating other source files, please choose filenames that are a complete single word: don't abbreviate (abbreviations are cryptic). For a minimal program containing all these files, see Section 11.4.2 [The TEMPLATE program], page 329.

`main.c` Each executable has a `main` function, which is located in `main.c`. Therefore this file in any program's source code will be the starting point. No actual processing functions must be defined in this file, the function(s) in this file are only meant to connect the most high level steps of each program. Generally, `main` will first call the top user interface function to read user input and make all the preparations. Then it will pass control to the top processing function for that program. The functions to do both these jobs must be defined in other source files.

`main.h` All the major parameters which will be used in the program must be stored in a structure which is defined in `main.h`. The name of this structure is usually `prognameparams`, for example `cropparams` or `noisechiselparams`. So `#include "main.h"` will be a staple in all the source codes of the program. It is also regularly the first (and only) argument most of the program's functions which greatly helps in readability.

Keeping all the major parameters of a program in this structure has the major benefit that most functions will only need one argument: a pointer to this structure. This will significantly facilitate the job of the programmer, the inquirer and the computer. All the programs in Gnuastro are designed to be low-level, small and independent parts, so this structure should not get too large.

The main root structure of all programs contains atleast one instance of the `gal_options_common_params` structure. This structure will keep the values to all common options in Gnuastro's programs (see Section 4.1.2 [Common options], page 52). This top root structure is conveniently called `p` (short for parameters) by all the functions in the programs and the common options parameters within it are called `cp`. With this convention any reader can immediately understand where to look for the definition of one parameter. For example you know that `p->cp->output` is in the common parameters while `p->threshold` is in the program's parameters.

With this basic root structure, source code of functions can potentially become full of structure de-reference operators (`->`) which can make the code very unreadable. In order to avoid this, whenever a structure element is used more than a couple of times in a function, a variable of the same type and with the

same name (so it can be searched) as the desired structure element should be defined with the value of the root structure inside of it in definition time. Here is an example.

```
char *hdu=p->cp.hdu;
float threshold=p->threshold;
```

`args.h` The options particular to each program are defined in this file. Each option is defined by a block of parameters in `program_options`. These blocks are all you should modify in this file, leave the bottom group of definitions untouched. These are fed directly into the GNU C library's Argp facilities and it is recommended to have a look at that for better understand what is going on, although this is not required here.

Each element of the block defining an option is described under `argp_option` in `bootstrapped/lib/argp.h` (from Gnuastro's top source file). Note that the last few elements of this structure are Gnuastro additions (not documented in the standard Argp manual). The values to these last elements are defined in `lib/gnuastro/type.h` and `lib/gnuastro-internal/options.h` (from Gnuastro's top source directory).

`ui.h` Besides declaring the exported functions of `ui.c`, this header also keeps the "key"s to every program-specific option. The first class of keys for the options that have a short-option version (single letter, see Section 4.1.1.2 [Options], page 50). The character that is defined here is the option's short option name. The list of available alphabet characters can be seen in the comments. Recall that some common options also take some characters, for those, see `lib/gnuastro-internal/options.h`.

The second group of options are those that don't have a short option alternative. Only the first in this group needs a value (`1000`), the rest will be given a value by C's `enum` definition, so the actual value is irrelevant and must never be used, always use the name.

`ui.c` Everything related to reading the user input arguments and options, checking the configuration files and checking the consistency of the input parameters before the actual processing is run should be done in this file. Since most functions are the same, with only the internal checks and structure parameters differing. We recommend going through the `ui.c` of Section 11.4.2 [The TEMPLATE program], page 329, or several other programs for a better understanding.

The most high-level function in `ui.c` is named `ui_read_check_inputs_setup`. It accepts the raw command-line inputs and a pointer to the root structure for that program (see the explanation for `main.h`). This is the function that `main` calls. The basic idea of the functions in this file is that the processing functions should need a minimum number of such checks. With this convention an inquirer who only wants to understand only one part (mostly the processing part and not user input details and sanity checks) of the code can easily do so in the later files. It also makes all the errors related to input appear before the processing begins which is more convenient for the user.

`progname.c, progname.h`

The high-level processing functions in each program are in a file named `progname.c`, for example `crop.c` or `noisechisel.c`. The function within these files which `main` calls is also named after the program, for example

```
void
crop(struct cropparams *p)
```

or

```
void
noisechisel(struct noisechiselparams *p)
```

In this manner, if an inquirer is interested the processing steps, they can immediately come and check this file for the first processing step without having to go through `main.c` and `ui.c` first. In most situations, any failure in any step of the programs will result in an informative error message and an immediate abort in the program. So there is usually no need for return values. Under more complicated situations where a return value might be necessary, `void` will be replaced with an `int` in the examples above. This value must be directly returned by `main`, so it has to be an `int`.

`authors-cite.h`

This header file keeps the global variable for the program authors and its BibTeX record for citation. They are used in the outputs of the common options `--version` and `--cite`, see Section 4.1.2.3 [Operating mode options], page 55.

11.4.2 The TEMPLATE program

The extra creativity offered by libraries comes at a cost: you have to actually write your `main` function and get your hands dirty in managing user inputs: are all the necessary parameters given a value? is the input in the correct format? do the options and the inputs correspond? and many other similar checks. So when an operation has well-defined inputs and outputs and is commonly needed, it is much more worthwhile to simply do use all the great features that Gnuastro has already defined for such operations.

To make it easier to learn/apply the internal program infra-structure discussed in Section 11.4.1 [Mandatory source code files], page 327, Gnuastro ships with a template program when using the Section 3.2.2 [Version controlled source], page 32. It is not available in the Gnuastro tarball so it doesn't confuse people using the tarball. The `bin/TEMPLATE` directory in Gnuastro's clone contains the bare-minimum files necessary to define a new program and all the necessary files/functions are pre-defined there. You can take the following steps if you want to add a new program to Gnuastro:

1. Select a name for your new program (for example `myprog`).
2. Copy the `TEMPLATE` directory to a directory with your program's name:

   ```
   $ cp -R bin/TEMPLATE bin/myprog
   ```

3. Open `configure.ac` in the top Gnuastro source. This file manages the operations that are done when a user runs `./configure`. Going down the file, you will notice repetitive parts for each program. Copy one of those and correct the names of the copied program to your new program name. We follow alphabetic ordering here, so please place it correctly. There are multiple places where this has to be done, so be

patient and go down to the bottom of the file. Ultimately add `bin/myprog/Makefile` to `AC_CONFIG_FILES`, only here the ordering depends on the length of the name.

4. Open `Makefile.am` in the top Gnuastro source. Similar to the previous step, add your new program similar to all the other programs.
5. Change `TEMPLATE` to `myprog` in the file names and contents of the files in the `bin/myprog/` directory.
6. Run the following command to re-build the configuration and build system.

```
$ autoreconf -f
```

Your new program will be built the next time you run `./configure` and `make`. You can now start adding your special checks/processing.

11.5 Documentation

Documentation (this book) is an integral part of Gnuastro (see Section 1.2 [Science and its tools], page 2). Documentation is not considered a separate project and must be written by its developers. Users can make edits/corrections, but the initial writing must be by the developer. So, no change is considered valid for implementation unless the respective parts of the book have also been updated. The following procedure can be a good suggestion to take when you have a new idea and are about to start implementing it.

The steps below are not a requirement, the important thing is that when you send the program to be included in Gnuastro, the book and the code have to both be fully up-to-date and compatible and the purpose should be very clearly explained. You can follow any path you choose to do this, the following path was what we have found to be most successful until now.

1. Edit the book and fully explain your desired change, such that your idea is completely embedded in the general context of the book with no sense of discontinuity for a first time reader. This will allow you to plan the idea much more accurately and in the general context of Gnuastro or a particular program. Later on, when you are coding, this general context will significantly help you as a road-map.

   A very important part of this process is the program introduction, which explains the purposes of the program. Before actually starting to code, explain your idea's purpose thoroughly in the start of the program section you wish to add or edit. While actually writing its purpose for a new reader, you will probably get some very valuable ideas that you hadn't thought of before. This has occurred several times during the creation of Gnuastro. If an introduction already exists, embed or blend your idea's purpose with the existing purposes. We emphasize that doing this is equally useful for you (as the programmer) as it is useful for the user (reader). Recall that the purpose of a program is very important, see Section 11.2 [Program design philosophy], page 322.

   As you have already noticed for every program, it is very important that the basics of the science and technique be explained in separate subsections prior to the 'Invoking Programname' subsection. If you are writing a new program or your addition to an existing program involves a new concept, also include such subsections and explain the concepts so a person completely unfamiliar with the concepts can get a general initial understanding. You don't have to go deep into the details, just enough to get an interested person (with absolutely no background) started. If you feel you can't

do that, then you have probably not understood the concept yourself! If you feel you don't have the time, then think about yourself as the reader in one year: you will forget almost all the details, so now that you have done all the theoretical preparations, add a few more hours and document it, so next time you don't have to prepare as much. Have in mind that your only limitation in length is the fatigue of the reader after reading a long text, nothing else. So as long as you keep it relevant/interesting for the reader, there is no page number limit/cost!

It might also help if you start discussing the usage of your idea in the 'Invoking ProgramName' subsection (explaining the options and arguments you have in mind) at this stage too. Actually starting to write it here will really help you later when you are coding.

2. After you have finished adding your initial intended plan to the book, then start coding your change or new program within the Gnuastro source files. While you are coding, you will notice that somethings should be different from what you wrote in the book (your initial plan). So correct them as you are actually coding, but don't worry too much about missing a few things (see the next step).
3. After your work has been fully implemented, read the section documentation from teh start and see if you didn't miss any change in the coding and to see if the context is fairly continuous for a first time reader (who hasn't seen the book or had known Gnuastro before you made your change).

11.6 Building and debugging

To build the various programs and libraries in Gnuastro, the GNU build system is used which defines the steps in Section 1.1 [Quick start], page 1. It consists of GNU Autoconf, GNU Automake and GNU Libtool which are collectively known as GNU Autotools. They provide a very portable system to check the hosts environment and compile Gnuastro based on that. They also make installing everything in their standard places very easy for the programmer. Most of the small caps files that you see in the top source directory of the tarball are created by these three tools (see Section 3.2.2 [Version controlled source], page 32). To facilitate the building and testing of your work during development, Gnuastro comes with two useful scripts:

`tmpfs-config-make`
: This is more fully described in Section 3.3.1.4 [Configure and build in RAM], page 44. During development, you will usually run this command only once (at the start of your work).

`tests/during-dev.sh`
: This script is designed to be run each time you make a change and want to test your work (with some possible input and output). The script itself is heavily commented and thoroughly describes the best way to use it, so we won't repeat it here.

  As a short summary: you specify the build directory, an output directory (for the built program to be run in, and also contains the inputs), the program's short name and the arguments and options that it should be run with. This script will then build Gnuastro, go to the output directory and run the built

executable from there. One option for the output directory might be your desktop, so you can easily see the output files and delete them when you are finished. The main purpose of these scripts is to keep your source directory clean and facilitate your development.

By default all the programs are compiled with optimization flags for increased speed. A side effect of optimization is that valuable debugging information is lost. All the libraries are also linked as shared libraries by default. Shared libraries further complicate the debugging process and significantly slow down the compilation (the `make` command). So during development it is recommended to configure Gnuastro as follows:

```
$ ./configure --disable-shared CFLAGS="-g -O0"
```

These options to configure are already included in `tmpfs-config-make`, you just have to un-comment them when you want to start developing or debugging.

In order to understand the building process, you can go through the Autoconf, Automake and Libtool manuals, like all GNU manuals they provide both a great tutorial and technical documentation. The "A small Hello World" section in Automake's manual (in chapter 2) can be a good starting guide after you have read the separate introductions.

11.7 Test scripts

As explained in Section 3.3.2 [Tests], page 45, for every program some simple tests are written to check the various independent features of the program. All the tests are placed in the `tests/` directory. The `tests/prepconf.sh` script is the first 'test' that will be run. It will copy all the configuration files from the various directories to a `tests/.gnuastro` directory (which it will make) so the various tests can set the default values. This script will also make sure the programs don't go searching for user and system wide configuration files to avoid the mixing of values with different Gnuastro version on the system.

For each program, the tests are placed inside directories with the program name. Each test is written as a shell script. The last line of this script is the test which runs the program with certain parameters. The return value of this script determines the fate of the test, see the "Support for test suites" chapter of the Automake manual for a very nice and complete explanation. In every script, two variables are defined at first: `prog` and `execname`. The first specifies the program name and the second the location of the executable.

The most important thing to have in mind about all the test scripts is that they are run from inside the `tests/` directory in the "build tree". Which can be different from the directory they are stored in (known as the "source tree"). The `tmpfs-config-make` script uses this feature (see Section 3.3.1.4 [Configure and build in RAM], page 44). This distinction is made by GNU Autoconf and Automake (which configure, build and install Gnuastro) so that you can install the program even if you don't have write access to the directory keeping the source files. See the "Parallel build trees (a.k.a VPATH builds)" in the Automake manual for a nice explanation.

Because of this, any necessary inputs that are distributed in the tarball[17], for example the catalogs necessary for checks in MakeProfiles and Crop, must be identified with the `$topsrc` prefix instead of `../` (for the top source directory that is unpacked). This `$topsrc`

---

[17] In many cases, the inputs of a test are outputs of previous tests, this doesn't apply to this class of inputs. Because all outputs of previous tests are in the "build tree".

variable points to the source tree where the script can find the source data (it is defined in `tests/Makefile.am`). The executables and other test products were built in the build tree (where they are being run), so they don't need to be prefixed with that variable. This is also true for images or files that were produced by other tests.

11.8 Developer's checklist

This is a checklist of things to do after applying your changes/additions in Gnuastro:

1. If the change is non-trivial, write test(s) in the `tests/progname/` directory to test the change(s)/addition(s) you have made. Then add their file names to `tests/Makefile.am`.
2. Run `$ make check` to make sure everything is working correctly.
3. Make sure the documentation (this book) is completely up to date with your changes, see Section 11.5 [Documentation], page 330.
4. Commit the change to your issue branch (see Section 11.11.3 [Production workflow], page 338, and Section 11.11.4 [Forking tutorial], page 339). Afterwards, run Autoreconf to generate the appropriate version number:

   ```
   $ autoreconf -f
   ```

5. Finally, to make sure everything will be built, installed and checked correctly run (after re-configuring, and re-building). To greatly speed up the process, use multiple threads (8 in the example below, change it appropriately)

   ```
   $ make distcheck -j8
   ```

   This command will create a distribution file (ending with `.tar.gz`) and try to compile it in the most general cases, then it will run the tests on what it has built in its own mini-environment. If `$ make distcheck` finishes successfully, then you are safe to send your changes to us to implement or for your own purposes. See Section 11.11.3 [Production workflow], page 338, and Section 11.11.4 [Forking tutorial], page 339.

11.9 Gnuastro project webpage

Gnuastro's central management hub (`https://savannah.gnu.org/projects/gnuastro/`)[18] is located on GNU Savannah (`https://savannah.gnu.org/`)[19]. Savannah is the central software development management system for many GNU projects. Through this central hub, you can view the list of activities that the developers are engaged in, their activity on the version controlled source, and other things. Each defined activity in the development cycle is known as an 'issue' (or 'item'). An issue can be a bug (see Section 1.7 [Report a bug], page 9), or a suggested feature (see Section 1.8 [Suggest new feature], page 11) or an enhancement or generally any *one* job that is to be done. In Savannah, issues are classified into three categories or 'tracker's:

Support
: This tracker is a way that (possibly anonymous) users can get in touch with the Gnuastro developers. It is a complement to the bug-gnuastro mailing list (see Section 1.7 [Report a bug], page 9). Anyone can post an issue to this tracker. The developers will not submit an issue to this list. They will only reassign the

[18] `https://savannah.gnu.org/projects/gnuastro/`

[19] `https://savannah.gnu.org/`

issues in this list to the other two trackers if they are valid[20]. Ideally (when the developers have time to put on Gnuastro, please don't forget that Gnuastro is a volunteer effort), there should be no open items in this tracker.

Bugs
: This tracker contains all the known bugs in Gnuastro (problems with the existing tools).

Tasks
: The items in this tracker contain the future plans (or new features/capabilities) that are to be added to Gnuastro.

All the trackers can be browsed by a (possibly anonymous) visitor, but to edit and comment on the Bugs and Tasks trackers, you have to be a registered on Savannah. When posting an issue to a tracker, it is very important to choose the 'Category' and 'Item Group' options accurately. The first contains a list of all Gnuastro's programs along with 'Installation', 'New program' and 'Webpage'. The "Item Group" contains the nature of the issue, for example if it is a 'Crash' in the software (a bug), or a problem in the documentation (also a bug) or a feature request or an enhancement.

The set of horizontal links on the top of the page (Starting with 'Main' and 'Homepage' and finishing with 'News') are the easiest way to access these trackers (and other major aspects of the project) from any part of the project webpage. Hovering your mouse over them will open a drop down menu that will link you to the different things you can do on each tracker (for example, 'Submit new' or 'Browse'). When you browse each tracker, you can use the "Display Criteria" link above the list to limit the displayed issues to what you are interested in. The 'Category' and 'Group Item' (explained above) are a good starting point.

Any new issue that is submitted to any of the trackers, or any comments that are posted for an issue, is directly forwarded to the gnuastro-devel mailing list (`https://lists.gnu.org/mailman/listinfo/gnuastro-devel`, see Section 11.10 [Developing mailing lists], page 335, for more). This will allow anyone interested to be up to date on the overall development activity in Gnuastro and will also provide an alternative (to Savannah) archiving for the development discussions. Therefore, it is not recommended to directly post an email to this mailing list, but do all the activities (for example add new issues, or comment on existing ones) on Savannah.

Do I need to be a member in Savannah to contribute to Gnuastro? No.

The full version controlled history of Gnuastro is available for anonymous download or cloning. See Section 11.11.3 [Production workflow], page 338, for a description of Gnuastro's Integration-Manager Workflow. In short, you can either send in patches, or make your own fork. If you choose the latter, you can push your changes to your own fork and inform us. We will then pull your changes and merge them into the main project. Please see Section 11.11.4 [Forking tutorial], page 339, for a tutorial.

[20] Some of the issues registered here might be due to a mistake on the user's side, not an actual bug in the program.

11.10 Developing mailing lists

To keep the developers and interested users up to date with the activity and discussions within Gnuastro, there are two mailing lists which you can subscribe to:

`gnuastro-devel@gnu.org`
(at `https://lists.gnu.org/mailman/listinfo/gnuastro-devel`)

: All the posts made in the support, bugs and tasks discussions of Section 11.9 [Gnuastro project webpage], page 333, are also sent to this mailing address and archived. By subscribing to this list you can stay up to date with the discussions that are going on between the developers before, during and (possibly) after working on an issue. All discussions are either in the context of bugs or tasks which are done on Savannah and circulated to all interested people through this mailing list. Therefore it is not recommended to post anything directly to this mailing list. Any mail that is sent to it from Savannah to this list has a link under the title "Reply to this item at:". That link will take you directly to the issue discussion page, where you can read the discussion history or join it.

  While you are posting comments on the Savannah issues, be sure to update the meta-data. For example if the task/bug is not assigned to anyone and you would like to take it, change the "Assigned to" box, or if you want to report that it has been applied, change the status and so on. All these changes will also be circulated with the email very clearly.

`gnuastro-commits@gnu.org`
(at `https://lists.gnu.org/mailman/listinfo/gnuastro-commits`)

: This mailing list is defined to circulate all commits that are done in Gnuastro's version controlled source, see Section 3.2.2 [Version controlled source], page 32. If you have any ideas, or suggestions on the commits, please use the bug and task trackers on Savannah to followup the discussion, do not post to this list. All the commits that are made for an already defined issue or task will state the respective ID so you can find it easily.

11.11 Contributing to Gnuastro

You have this great idea or have found a good fix to a problem which you would like to implement in Gnuastro. You have also become familiar with the general design of Gnuastro in the previous sections of this chapter (see Chapter 11 [Developing], page 320) and want to start working on and sharing your new addition/change with the whole community as part of the official release. This is great and your contribution is most welcome. This section and the next (see Section 11.8 [Developer's checklist], page 333) are written in the hope of making it as easy as possible for you to share your great idea with the Gnuastro community.

In this section we discuss the final steps you have to take: legal and technical. From the legal perspective, the copyright of any work you do on Gnuastro has to be assigned to the Free Software Foundation (FSF) and the GNU operating system, or you have to sign a disclaimer. We do this to ensure that Gnuastro can remain free in the future, see Section 11.11.1 [Copyright assignment], page 336. From the technical point of view, in this section we also discuss commit guidelines (Section 11.11.2 [Commit guidelines], page 337) and the general version control workflow of Gnuastro in Section 11.11.3 [Production workflow], page 338, along with a tutorial in Section 11.11.4 [Forking tutorial], page 339.

Recall that before starting the work on your idea, be sure to checkout the bugs and tasks trackers in Section 11.9 [Gnuastro project webpage], page 333, and announce your work there so you don't end up spending time on something others have already worked on, and also to attract similarly interested developers to help you.

11.11.1 Copyright assignment

Gnuastro's copyright is owned by the FSF. Professor Eben Moglen, of the Columbia University Law School has given a nice summary of the reasons for this at `https://www.gnu.org/licenses/why-assign`. Below we are copying it verbatim for self consistency (in case you are offline or reading in print).

> Under US copyright law, which is the law under which most free software programs have historically been first published, there are very substantial procedural advantages to registration of copyright. And despite the broad right of distribution conveyed by the GPL, enforcement of copyright is generally not possible for distributors: only the copyright holder or someone having assignment of the copyright can enforce the license. If there are multiple authors of a copyrighted work, successful enforcement depends on having the cooperation of all authors.
>
> In order to make sure that all of our copyrights can meet the record keeping and other requirements of registration, and in order to be able to enforce the GPL most effectively, FSF requires that each author of code incorporated in FSF projects provide a copyright assignment, and, where appropriate, a disclaimer of any work-for-hire ownership claims by the programmer's employer. That way we can be sure that all the code in FSF projects is free code, whose freedom we can most effectively protect, and therefore on which other developers can completely rely.

Please get in touch with the Gnuastro maintainer (currently Mohammad Akhlaghi, akhlaghi -at- gnu -dot- org) to follow the procedures. It is possible to do this for each change (good for for a single contribution), and also more generally for all the changes/additions you do in the future within Gnuastro. So if you have already assigned the copyright of your work on another GNU software to the FSF, it should be done again for Gnuastro. The FSF has staff working on these legal issues and the maintainer will get you in touch with them to do the paperwork. The maintainer will just be informed in the end so your contributions can be merged within the Gnuastro source code.

Gnuastro will gratefully acknowledge (see Section 1.11 [Acknowledgments], page 12) all the people who have assigned their copyright to the FSF and have thus helped to guarantee the freedom and reliability of Gnuastro. The Free Software Foundation will also acknowledge your copyright contributions in the Free Software Supporter: `https://www.fsf.org/free-software-supporter` which will circulate to a very large community (104,444 people in April 2016). See the archives for some examples and subscribe to receive interesting updates. The very active code contributors (or developers) will also be recognized as project members on the Gnuastro project webpage (see Section 11.9 [Gnuastro project webpage], page 333) and can be given a `gnu.org` email address. So your very valuable contribution and copyright assignment will not be forgotten and is highly appreciated by a very large community. If you are reluctant to sign an assignment, a disclaimer is also acceptable.

Do I need a disclaimer from my university or employer? It depends on the contract with your university or employer. From the FSF's `/gd/gnuorg/conditions.text`: "If you are employed to do programming, or have made an agreement with your employer that says it owns programs you write, we need a signed piece of paper from your employer disclaiming rights to" Gnuastro. The FSF's copyright clerk will kindly help you decide, please consult the following email address: "assign -at- gnu -dot- org".

11.11.2 Commit guidelines

To be able to cleanly integrate your work with the other developers, **never commit on the `master` branch** (see Section 11.11.3 [Production workflow], page 338, for a complete discussion and Section 11.11.4 [Forking tutorial], page 339, for a cookbook example). In short, leave `master` only for changes you fetch, or pull from the official repository (see Section 3.2.2.2 [Synchronizing], page 35).

In the Gnuastro commit messages, we strive to follow these standards. Note that in the early phases of Gnuastro's development, we are experimenting and so if you notice earlier commits don't satisfy some of the guidelines below, it is because they predate that guideline.

Commit title

The commits have to start with one short descriptive title. The title is separated from the body with one blank line. Run `git log` to see some of the most recent commit messages as an example. In general, the title should satisfy the following conditions:

- It is best for the title to be short, about 60 (or even 50) characters. Most emulated command-line terminals are about 80 characters wide. However, we should also allow for the commit hashes which are printed in `git log --oneline`, and also branch names or the graph structure outputs of `git log` which are also commonly used.
- The title should not finish with any full-stops or periods ('`.`').

Commit body

The body of the commit message is separated from the title by one empty line. Recall that anyone who has subscribed to `gnuastro-commits` mailing list will get the commit in their email after it has been pushed to `master`. People will also read them when they synchronize with the main Gnuastro repository (see Section 3.2.2.2 [Synchronizing], page 35). Finally, the commit messages will later be used to update the `NEWS` file on each release. Therefore the commit message body plays a very important role in the development of Gnuastro, so please adhere to the following guidelines.

- The body should be very descriptive. Start the commit message body by explaining what changes your commit makes from a user's perspective (added, changed, or removed options, or arguments to programs or libraries, or modified algorithms, or new installation step, or etc).
- Try to explain the committed contents as best as you can. Recall that the readers of your commit message do not necessarily have your current background. After some time you will also forget the context, so this request is not just for others[21]. Therefore be very descriptive and explain

[21] `http://catb.org/esr/writings/unix-koans/prodigy.html`

as much as possible: what the bug/task was, justify the way you fixed it and discuss other possible solutions that you might not have included. For the last item, it is best to discuss them thoroughly as comments in the appropriate section of the code, but only give a short summary in the commit message. Note that all added and removed source code lines will also be circulated in the `gnuastro-commits` mailing list.

- Like all other Gnuastro's text files, the lines in the commit body should not be longer than 75 characters, see Section 11.3 [Coding conventions], page 323. This is to ensure that on standard terminal emulators (with 80 character width), the `git log` output can be cleanly displayed (note that the commit message is indented in the output of `git log`). If you use Emacs, Gnuastro's `.dir-locals.el` file will ensure that your commits satisfy this condition (using `M-q`).
- When the commit is related to a task or a bug, please include the respective ID (in the format of `bug/task #ID`, note the space) in the commit message (from Section 11.9 [Gnuastro project webpage], page 333) for interested people to be able to followup the discussion that took place there. If the commit fixes a bug or finishes a task, the recommended way is to add a line after the body with '`This fixes bug #ID.`', or '`This finishes task #ID.`'. Don't assume that the reader has internet access to check the bug's full description when reading the commit message, so give a short introduction too.

11.11.3 Production workflow

Fortunately 'Pro Git' has done a wonderful job in explaining the different workflows in Chapter 5[22] and in particular the "Integration-Manager Workflow" explained there. The implementation of this workflow is nicely explained in Section 5.2[23] under "Forked-Public-Project". We have also prepared a short tutorial in Section 11.11.4 [Forking tutorial], page 339. Anything on the master branch should always be tested and ready to be built and used. As described in 'Pro Git', there are two methods for you to contribute to Gnuastro in the Integration-Manager Workflow:

1. You can send commit patches by email as fully explained in 'Pro Git'. This is good for your first few contributions. Just note that raw patches (containing only the diff) do not have any meta-data (author name, date and etc). Therefore they will not allow us to fully acknowledge your contributions as an author in Gnuastro: in the `AUTHORS` file and at the start of the PDF book. These author lists are created automatically from the version controlled source.

   To receive full acknowledgement when submitting a patch, is thus advised to use Git's `format-patch` tool. See Pro Git's Public project over email (`https://git-scm.com/book/en/v2/Distributed-Git-Contributing-to-a-Project#Public-Project-over-Email`) section for a nice explanation. If you would like to get more heavily involved in Gnuastro's development, then you can try the next solution.

[22] `http://git-scm.com/book/en/v2/Distributed-Git-Distributed-Workflows`
[23] `http://git-scm.com/book/en/v2/Distributed-Git-Contributing-to-a-Project`

2. You can have your own forked copy of Gnuastro on any hosting site you like (GitHub, GitLab, BitBucket, or etc) and inform us when your changes are ready so we merge them in Gnuastro. This is more suited for people who commonly contribute to the code (see Section 11.11.4 [Forking tutorial], page 339).

In both cases, your commits (with your name and information) will be preserved and your contributions will thus be fully recorded in the history of Gnuastro and in the `AUTHORS` file and this book (second page in the PDF format) once they have been incorporated into the official repository. Needless to say that in such cases, be sure to follow the bug or task trackers (or subscribe to the `gnuastro-devel` mailing list) and contact us before hand so you don't do something that someone else is already working on. In that case, you can get in touch with them and help the job go on faster, see Section 11.9 [Gnuastro project webpage], page 333. This workflow is currently mostly borrowed from the general recommendations of Git[24] and GitHub. But since Gnuastro is currently under heavy development, these might change and evolve to better suit our needs.

11.11.4 Forking tutorial

This is a tutorial on the second suggested method that you can submit your modifications in Gnuastro (see Section 11.11.3 [Production workflow], page 338), which is commonly known as forking. Let's assume you want to start contributing to Gnuastro and you would like to use GitLab[25] to host your work remotely and share it with other Gnuastro developers once you are ready. You can create an empty repository on your hosting service webpage. If this is your first hosted repository on the webpage, you also have to upload your public SSH key[26] for the `git push` command below to work. Here we'll assume you use the name `janedoe` to refer to yourself everywhere and that you choose `gnuastro-janedoe` as the name of your Gnuastro fork. Any online hosting service will give you an address (similar to '`git@gitlab.com:`' below) of the empty repository you have created using their webpage, use that address in the third line below.

```
$ git clone git://git.sv.gnu.org/gnuastro.git
$ cd gnuastro
$ git remote add janedoe git@gitlab.com:janedoe/gnuastro-janedoe.git
$ git push janedoe master
```

The full Gnuastro history is now pushed onto your hosting service and the `janedoe` remote is now also following your `master` branch. If you run `git remote show REMOTENAME` for the `origin` and `janedoe` remotes, you will see their difference: the first has pull access and the second doesn't. This nicely summarizes the main idea behind this workflow: you push to your remote repository, we pull from it and merge it into `master`, then you finalize it by pulling from the main repository.

To test (compile) your changes during your work, you will need to bootstrap the version controlled source, see Section 3.2.2.1 [Bootstrapping], page 33, for a full description. The process above is only necessary for your first time setup, you don't need to repeat it. However, please repeat the steps below for each independent issue you intend to work on.

---

[24] `https://github.com/git/git/blob/master/Documentation/SubmittingPatches`

[25] See `https://www.gnu.org/software/repo-criteria-evaluation.html` for an evaluation of the major existing repositories.

[26] For example see this explanation provided by GitLab: `http://docs.gitlab.com/ce/ssh/README.html`.

Let's assume you have found a bug in `lib/statistics.c`'s median calculating function. You announce it (see Section 1.7 [Report a bug], page 9) so everyone knows you are working on it. However, if it is a very simple, clear/obvious and straightforward bug to fix you can skip this initial announcement. With the commands below, you make a branch, checkout to it, correct the bug, check if it is indeed fixed, add it to the staging area, commit it to the new branch and push it to your hosting service. But before all of them, make sure that you are on the `master` branch and that your `master` branch is up to date with the main Gnuastro repository with the first two commands.

```
$ git checkout master
$ git pull
$ git checkout -b bug-median-stats      # Choose a descriptive name
$ emacs lib/statistics.c
$                                       # do your checks here
$ git add lib/statistics.c
$ git commit
$ git push janedoe bug-median-stats
```

Your new branch is now on your hosted repository. Through the respective tacker on Savannah (see Section 11.9 [Gnuastro project webpage], page 333) you can then let the other developers know that your `bug-median-stats` branch is ready. They will pull your work, test it themselves and if it is ready to be merged into the main Gnuastro history, they will merge it into the `master` branch. After that is done, you can simply checkout your local `master` branch and pull all the changes from the main repository. After the pull you can run '`git log`' as shown below, to see how `bug-median-stats` is merged with master. So you can push all the changes to your hosted repository and delete the branches:

```
$ git checkout master
$ git pull
$ git log --oneline --graph --decorate --all
$ git push janedoe master
$ git branch -d bug-median-stats                # delete local branch
$ git push janedoe --delete bug-median-stats    # delete remote branch
```

Just as a reminder, always keep your work on each issue in a separate local and remote branch so work can progress on them independently. After you make your announcement, other people might contribute to the branch before merging it in to `master`, so this is very important. Also before starting each issue branch from `master`, be sure to run `git pull` in `master` as shown above, to start your branch (work) from the most recent history point and thus simplify the final merging of your work.

Appendix A Gnuastro programs list

GNU Astronomy Utilities 0.3, contains the following programs. They are sorted in alphabetical order and followed by their version number. A short description is provided for each program which starts with the executable names in `thisfont` followed by a link to the respective section in parenthesis, see Section 1.4 [Naming convention], page 5. Throughout this book, they are ordered based on their context, please see the book contents for contextual ordering.

Arithmetic
: (`astarithmetic`, see Section 6.2 [Arithmetic], page 108) For arithmetic operations on multiple (theoretically unlimited) number of datasets (images). It has a large and growing set of arithmetic, mathematical, and even statistical operators (for example `+`, `-`, `*`, `/`, `sqrt`, `log`, `min`, `average`, `median`).

BuildProgram
: (`astbuildprog`, see Section 10.2 [BuildProgram], page 230) Compile, link and run programs that depend on the Gnuastro library (see Section 10.3 [Gnuastro library], page 233). This program will automatically link with the libraries that Gnuastro depends on, so there is no need to explicily mention them every time you are compiling a Gnuastro library dependent program.

ConvertType
: (`astconvertt`, see Section 5.2 [ConvertType], page 87) Convert astronomical data files (FITS or IMH) to and from several other standard image and data formats, for example TXT, JPEG, EPS or PDF.

Convolve
: (`astconvolve`, see Section 6.3 [Convolve], page 117) Convolve (blur or smooth) data with a given kernel in spatial and frequency domain on multiple threads. Convolve can also do de-convolution to find the appropriate kernel to PSF-match two images.

CosmicCalculator
: (`astcosmiccal`, see Section 9.1 [CosmicCalculator], page 216) Do cosmological calculations, for example the luminosity distance, distance modulus, comoving volume and many more.

Crop
: (`astcrop`, see Section 6.1 [Crop], page 97) Crop region(s) from an image and stitch several images if necessary. Inputs can be in pixel coordinates or world coordinates.

Fits
: (`astfits`, see Section 5.1 [Fits], page 80) View and manipulate FITS file extensions and header keywords.

Statistics
: (`aststatistics`, see Section 7.1 [Statistics], page 148) Get pixel statistics and save histogram and cumulative frequency plots.

MakeCatalog
: (`astmkcatalog`, see Section 7.3 [MakeCatalog], page 175) Make catalog of labeled image (output of NoiseChisel). The catalogs are highly customizable and adding new calculations/columns is very straightforward.

MakeNoise
: (`astmknoise`, see Section 8.2 [MakeNoise], page 210) Make (add) noise to an image, with a large set of random number generators and any seed.

MakeProfiles
: (`astmkprof`, see Section 8.1 [MakeProfiles], page 195) Make mock 2D profiles in an image. The central regions of radial profiles are made with a configurable 2D Monte Carlo integration. It can also build the profiles on an over-sampled image.

NoiseChisel
: (`astnoisechisel`, see Section 7.2 [NoiseChisel], page 163) Detect and segment signal in noise. It uses a technique to detect very faint and diffuse, irregularly shaped signal in noise (galaxies in the sky), using thresholds that are below the Sky value, see arXiv:1505.01664 (`http://arxiv.org/abs/1505.01664`).

Table
: (`asttable`, Section 5.3 [Table], page 94) Convert FITS binary and ASCII tables into other such tables, print them on the command-line, save them in a plain text file, or get the FITS table information.

Warp
: (`astwarp`, see Section 6.4 [Warp], page 138) Warp image to new pixel grid. Any projective transformation or Homography can be applied to the input images.

Appendix B Other useful software

In this appendix the installation of programs and libraries that are not direct Gnuastro dependencies are discussed. However they can be useful for working with Gnuastro.

B.1 SAO ds9

SAO ds9[1] is not a requirement of Gnuastro, it is a FITS image viewer. So to check your inputs and outputs, it is one of the best options. Like the other packages, it might already be available in your distribution's repositories. It is already pre-compiled in the download section of its webpage. Once you download it you can unpack and install (move it to a system recognized directory) with the following commands (`x.x.x` is the version number):

```
$ tar xf ds9.linux64.x.x.x.tar.gz
$ sudo mv ds9 /usr/local/bin
```

Once you run it, there might be a complaint about the Xss library, which you can find in your distribution package management system. You might also get an `XPA` related error. In this case, you have to add the following line to your `~/.bashrc` and `~/.profile` file (you will have to log out and back in again for the latter):

```
export XPA_METHOD=local
```

B.1.1 Viewing multiextension FITS images

The FITS definition allows for multiple extensions inside a FITS file, each extension can have a completely independent data set inside of it. If you ordinarily open a multi-extension FITS file with SAO ds9, for example by double clicking on the file or running `$ds9 foo.fits`, SAO ds9 will only show you the first extension. To be able to switch between the extensions you have to follow these menus in the SAO ds9 window: File→Open Other→Open Multi Ext Cube and then choose the Multi extension FITS file in your computer's file structure.

The method above is a little tedious to do every time you want view a multi-extension FITS file. Fortunately SAO ds9 also provides options that you can use to specify a particular behavior. One of those options is `-mecube` which opens a FITS image as a multi-extension data cube. So on the command-line, if you run `$ds9 -mecube foo.fits` a small window will also be opened, which allows you to switch between the image extensions that `foo.fits` might have. If `foo.fits` only consists of one extension, then SAO ds9 will open as usual.

Just to avoid confusion, note that SAO ds9 does not follow the GNU style of separating long and short options as explained in Section 4.1.1 [Arguments and options], page 48. In the GNU style, this 'long' option should have been called like `--mecube`, but SAO ds9 does follow those conventions and has its own.

It is really convenient if you set ds9 to always run with the `-mecube` option on your graphical display. On GNOME 3 (the most popular graphic user interface for GNU/Linux systems) you can do this by taking the following steps:

- Open your favorite text editor and put the following text in a file that ends with `.desktop`, for example `saods9.desktop`. The file is very descriptive.

  ```
  [Desktop Entry]
  ```

[1] `http://ds9.si.edu/`

```
Type=Application
Version=1.0
Name=SAO ds9
Comment=View FITS images
Exec=ds9 -mecube %f
Terminal=false
Categories=Graphic;FITS;
```

- Copy this file into your local (user) applications directory:

  ```
  $ cp saods9.desktop ~/.local/share/applications/
  ```

  In case you don't have the directory, you can make it yourself:

  ```
  $ mkdir -p ~/.local/share/applications/
  ```

- The steps above will add SAO ds9 as one of your applications. To make it default for every time you click on a FITS file. Right click on a FITS file and select "Open With", then go into "Other Application..." and choose "SAO ds9".

In case you are using GNOME 2 you can take the following steps: right click on a FITS file and choose Properties→Open With→Add button. A list of applications will show up, ds9 might already be present in the list, but don't choose it because it will run with no options. Below the list is an option "Use a custom command". Click on it and write the following command: `ds9 -mecube` in the box and click "Add". Then finally choose the command you just added as the default and click the "Close" button.

B.2 PGPLOT

PGPLOT is a package for making plots in C. It is not directly needed by Gnuastro, but can be used by WCSLIB, see Section 3.1.1.3 [WCSLIB], page 28. As explained in Section 3.1.1.3 [WCSLIB], page 28, you can install WCSLIB without it too. It is very old (the most recent version was released early 2001!), but remains one of the main packages for plotting directly in C. WCSLIB uses this package to make plots if you want it to make plots. If you are interested you can also use it for your own purposes.

If you want your plotting codes in between your C program, PGPLOT is currently one of your best options. The recommended alternative to this method is to get the raw data for the plots in text files and input them into any of the various more modern and capable plotting tools separately, for example the Matplotlib library in Python or PGFplots in LaTeX. This will also significantly help code readability. Let's get back to PGPLOT for the sake of WCSLIB. Installing it is a little tricky (mainly because it is so old!).

You can download the most recent version from the FTP link in its webpage[2]. You can unpack it with the `tar xf` command. Let's assume the directory you have unpacked it to is `PGPLOT`, most probably it is: `/home/username/Downloads/pgplot/`. open the `drivers.list` file:

```
$ gedit drivers.list
```

Remove the `!` for the following lines and save the file in the end:

```
PSDRIV 1 /PS
PSDRIV 2 /VPS
```

[2] `http://www.astro.caltech.edu/~tjp/pgplot/`

```
PSDRIV 3 /CPS
PSDRIV 4 /VCPS
XWDRIV 1 /XWINDOW
XWDRIV 2 /XSERVE
```

Don't choose GIF or VGIF, there is a problem in their codes.

Open the `PGPLOT/sys_linux/g77_gcc.conf` file:

```
$ gedit PGPLOT/sys_linux/g77_gcc.conf
```

change the line saying: `FCOMPL="g77"` to `FCOMPL="gfortran"`, and save it. This is a very important step during the compilation of the code if you are in GNU/Linux. You now have to create a folder in `/usr/local`, don't forget to replace `PGPLOT` with your unpacked address:

```
$ su
# mkdir /usr/local/pgplot
# cd /usr/local/pgplot
# cp PGPLOT/drivers.list ./
```

To make the Makefile, type the following command:

```
# PGPLOT/makemake PGPLOT linux g77_gcc
```

It should finish by saying: `Determining object file dependencies`. You have done the hard part! The rest is easy: run these three commands in order:

```
# make
# make clean
# make cpg
```

Finally you have to place the position of this directory you just made into the `LD_LIBRARY_PATH` environment variable and define the environment variable `PGPLOT_DIR`. To do that, you have to edit your `.bashrc` file:

```
$ cd ~
$ gedit .bashrc
```

Copy these lines into the text editor and save it:

```
PGPLOT_DIR="/usr/local/pgplot/"; export PGPLOT_DIR
LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/usr/local/pgplot/
export LD_LIBRARY_PATH
```

You need to log out and log back in again so these definitions take effect. After you logged back in, you want to see the result of all this labor, right? Tim Pearson has done that for you, create a temporary folder in your home directory and copy all the demonstration files in it:

```
$ cd ~
$ mkdir temp
$ cd temp
$ cp /usr/local/pgplot/pgdemo* ./
$ ls
```

You will see a lot of pgdemoXX files, where XX is a number. In order to execute them type the following command and drink your coffee while looking at all the beautiful plots! You are now ready to create your own.

```
$ ./pgdemoXX
```

Appendix C GNU Free Doc. License

Version 1.3, 3 November 2008

Copyright © 2000, 2001, 2002, 2007, 2008 Free Software Foundation, Inc.
http://fsf.org/

Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

0. PREAMBLE

   The purpose of this License is to make a manual, textbook, or other functional and useful document *free* in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondarily, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

   This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

   We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

   This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you". You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

   A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

   A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

   The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released

under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

The "publisher" means any person or entity that distributes copies of the Document to the public.

A section "Entitled XYZ" means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as "Acknowledgements", "Dedications", "Endorsements", or "History".) To "Preserve the Title" of such a section when you modify the Document means that it remains a section "Entitled XYZ" according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any,

be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.

B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
D. Preserve all the copyright notices of the Document.
E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
H. Include an unaltered copy of this License.
I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their

titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

   You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

   The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

   In the combination, you must combine any sections Entitled "History" in the various original documents, forming one section Entitled "History"; likewise combine any sections Entitled "Acknowledgements", and any sections Entitled "Dedications". You must delete all sections Entitled "Endorsements."

6. COLLECTIONS OF DOCUMENTS

   You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

   You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

   A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an "aggregate" if the copyright resulting from the compilation is not used to limit the legal rights of the compilation's users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

   If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document's Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. TRANSLATION

   Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

   If a section in the Document is Entitled "Acknowledgements", "Dedications", or "History", the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. TERMINATION

   You may not copy, modify, sublicense, or distribute the Document except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, or distribute it is void, and will automatically terminate your rights under this License.

   However, if you cease all violation of this License, then your license from a particular copyright holder is reinstated (a) provisionally, unless and until the copyright holder explicitly and finally terminates your license, and (b) permanently, if the copyright holder fails to notify you of the violation by some reasonable means prior to 60 days after the cessation.

   Moreover, your license from a particular copyright holder is reinstated permanently if the copyright holder notifies you of the violation by some reasonable means, this is the first time you have received notice of violation of this License (for any work) from that copyright holder, and you cure the violation prior to 30 days after your receipt of the notice.

   Termination of your rights under this section does not terminate the licenses of parties who have received copies or rights from you under this License. If your rights have been terminated and not permanently reinstated, receipt of a copy of some or all of the same material does not give you any rights to use it.

10. FUTURE REVISIONS OF THIS LICENSE

    The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See `http://www.gnu.org/copyleft/`.

    Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation. If the Document specifies that a proxy can decide which future versions of this License can be used, that proxy's public statement of acceptance of a version permanently authorizes you to choose that version for the Document.

11. RELICENSING

    "Massive Multiauthor Collaboration Site" (or "MMC Site") means any World Wide Web server that publishes copyrightable works and also provides prominent facilities for anybody to edit those works. A public wiki that anybody can edit is an example of such a server. A "Massive Multiauthor Collaboration" (or "MMC") contained in the site means any set of copyrightable works thus published on the MMC site.

    "CC-BY-SA" means the Creative Commons Attribution-Share Alike 3.0 license published by Creative Commons Corporation, a not-for-profit corporation with a principal place of business in San Francisco, California, as well as future copyleft versions of that license published by that same organization.

    "Incorporate" means to publish or republish a Document, in whole or in part, as part of another Document.

    An MMC is "eligible for relicensing" if it is licensed under this License, and if all works that were first published under this License somewhere other than this MMC, and subsequently incorporated in whole or in part into the MMC, (1) had no cover texts or invariant sections, and (2) were thus incorporated prior to November 1, 2008.

    The operator of an MMC Site may republish an MMC contained in the site under CC-BY-SA on the same site at any time before August 1, 2009, provided the MMC is eligible for relicensing.

ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

```
Copyright (C)  year  your name.
Permission is granted to copy, distribute and/or modify this document
under the terms of the GNU Free Documentation License, Version 1.3
or any later version published by the Free Software Foundation;
with no Invariant Sections, no Front-Cover Texts, and no Back-Cover
Texts.  A copy of the license is included in the section entitled ``GNU
Free Documentation License''.
```

If you have Invariant Sections, Front-Cover Texts and Back-Cover Texts, replace the "with...Texts." line with this:

```
with the Invariant Sections being list their titles, with
the Front-Cover Texts being list, and with the Back-Cover Texts
being list.
```

If you have Invariant Sections without Cover Texts, or some other combination of the three, merge those two alternatives to suit the situation.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

Index: Macros, structures and functions

All Gnuastro library's exported macros start with GAL_, and its exported structures and functions start with gal_. This abbreviation stands for *G*NU *A*stronomy *L*ibrary. The next element in the name is the name of the header which declares or defines them, so to use the gal_array_fset_const function, you have to #include <gnuastro/array.h>. See Section 10.3 [Gnuastro library], page 233, for more. The pthread_barrier constructs are our implementation and are only available on systems that don't have them, see Section 10.3.2.1 [Implementation of pthread_barrier], page 236.

G

gal_arithmetic 288
gal_binary_connected_adjacency_matrix 309
gal_binary_connected_components 308
gal_binary_dilate 308
gal_binary_erode 308
gal_binary_fill_holes 309
gal_binary_open 308
gal_blank_alloc_write 244
gal_blank_as_string 245
gal_blank_flag 245
gal_blank_initialize 244
gal_blank_present 244
gal_blank_remove 245
gal_blank_write 244
gal_box_border_from_center 298
gal_box_ellipse_in_box 298
gal_box_overlap 298
gal_convolve_spatial 310
gal_convolve_spatial_correct_ch_edge 310
gal_data_alloc 251
gal_data_array_calloc 252
gal_data_array_free 252
gal_data_calloc_array 251
gal_data_copy 252
gal_data_copy_string_to_number 253
gal_data_copy_to_allocated 253
gal_data_copy_to_new_type 253
gal_data_copy_to_new_type_free 253
gal_data_dsize_is_different 250
gal_data_free 251
gal_data_free_contents 251
gal_data_initialize 251
gal_data_malloc_array 250
gal_data_ptr_dist 250
gal_data_ptr_increment 250
gal_dimension_add_coords 254
gal_dimension_coord_to_index 254
gal_dimension_dist_manhattan 254
gal_dimension_increment 254
gal_dimension_index_to_coord 254
gal_dimension_num_neighbors 254
gal_dimension_total_size 254
gal_fits_bitpix_to_type 269
gal_fits_datatype_to_type 269
gal_fits_hdu_format 270
gal_fits_hdu_num 270
gal_fits_hdu_open 270
gal_fits_hdu_open_format 270
gal_fits_img_info 274
gal_fits_img_read 274
gal_fits_img_read_kernel 274
gal_fits_img_read_to_type 274
gal_fits_img_write 275
gal_fits_img_write_corr_wcs_str 275
gal_fits_img_write_to_ptr 275
gal_fits_img_write_to_type 275
gal_fits_io_error 268
gal_fits_key_clean_str_value 271
gal_fits_key_list_add 273
gal_fits_key_list_add_end 273
gal_fits_key_read 272
gal_fits_key_read_from_ptr 271
gal_fits_key_write 273
gal_fits_key_write_filename 273
gal_fits_key_write_version 273
gal_fits_key_write_wcsstr 273
gal_fits_name_is_fits 268
gal_fits_name_save_as_string 268
gal_fits_open_to_write 270
gal_fits_suffix_is_fits 268
gal_fits_tab_format 276
gal_fits_tab_info 276
gal_fits_tab_read 276
gal_fits_tab_size 276
gal_fits_tab_write 276
gal_fits_type_to_bin_tform 269
gal_fits_type_to_bitpix 269
gal_fits_type_to_datatype 269
gal_git_describe 311
gal_interpolate_close_neighbors 311
gal_list_data_add 267
gal_list_data_add_alloc 267
gal_list_data_free 267
gal_list_data_number 267
gal_list_data_pop 267
gal_list_data_reverse 267
gal_list_dosizet_add 266
gal_list_dosizet_free 267
gal_list_dosizet_pop_smallest 266

gal_list_dosizet_print 266
gal_list_dosizet_to_sizet 266
gal_list_f32_add 261
gal_list_f32_free 262
gal_list_f32_number 261
gal_list_f32_pop 261
gal_list_f32_print 262
gal_list_f32_reverse 262
gal_list_f32_to_array 262
gal_list_f64_add 263
gal_list_f64_free 263
gal_list_f64_number 263
gal_list_f64_pop 263
gal_list_f64_print 263
gal_list_f64_reverse 263
gal_list_f64_to_array 263
gal_list_i32_add 258
gal_list_i32_free 259
gal_list_i32_number 259
gal_list_i32_pop 258
gal_list_i32_print 259
gal_list_i32_reverse 259
gal_list_i32_to_array 259
gal_list_osizet_add 265
gal_list_osizet_pop 265
gal_list_osizet_to_sizet_free 265
gal_list_sizet_add 260
gal_list_sizet_free 261
gal_list_sizet_number 260
gal_list_sizet_pop 260
gal_list_sizet_print 260
gal_list_sizet_reverse 260
gal_list_sizet_to_array 261
gal_list_str_add 257
gal_list_str_free 258
gal_list_str_number 257
gal_list_str_pop 257
gal_list_str_print 257
gal_list_str_reverse 258
gal_list_void_add 264
gal_list_void_free 264
gal_list_void_number 264
gal_list_void_pop 264
gal_list_void_reverse 264
gal_permutation_apply 302
gal_permutation_apply_inverse 302
gal_permutation_check 302
gal_polygon_area 300
gal_polygon_clip 300
gal_polygon_ordered_corners 299
gal_polygon_pin 300
gal_polygon_ppropin 300
gal_qsort_index_arr 300
gal_qsort_index_float_decreasing 300
gal_qsort_TYPE_decreasing 301
gal_qsort_TYPE_increasing 301
gal_statistics_cfp 306
gal_statistics_histogram 306
gal_statistics_is_sorted 305
gal_statistics_maximum 303
gal_statistics_mean 303
gal_statistics_mean_std 303
gal_statistics_median 303
gal_statistics_minimum 303
gal_statistics_mode 304
gal_statistics_mode_mirror_plots 305
gal_statistics_no_blank_sorted 305
gal_statistics_number 303
gal_statistics_quantile 304
gal_statistics_quantile_function 304
gal_statistics_quantile_function_index 304
gal_statistics_quantile_index 304
gal_statistics_regular_bins 305
gal_statistics_sigma_clip 306
gal_statistics_sort_decreasing 305
gal_statistics_sort_increasing 305
gal_statistics_std 303
gal_statistics_sum 303
gal_table_comments_add_intro 283
gal_table_info 282
gal_table_print_info 283
gal_table_read 283
gal_table_write 283
gal_table_write_log 284
gal_threads_attr_barrier_init 238
gal_threads_dist_in_threads 238
gal_threads_number 237
gal_threads_spin_off 238
gal_tile_block 290
gal_tile_block_blank_flag 291
gal_tile_block_check_tiles 291
gal_tile_block_increment 290
gal_tile_block_relative_to_other 291
gal_tile_block_write_const_value 291
gal_tile_full 295
gal_tile_full_free_contents 298
gal_tile_full_id_from_coord 298
gal_tile_full_permutation 296
gal_tile_full_sanity_check 296
gal_tile_full_two_layers 296
gal_tile_full_values_smooth 297
gal_tile_full_values_write 297
gal_tile_series_from_minmax 290
gal_tile_start_coord 289
gal_tile_start_end_coord 289
gal_tile_start_end_ind_inclusive 290
gal_txt_image_read 280
gal_txt_line_stat 280
gal_txt_table_info 280
gal_txt_table_read 280
gal_txt_write 280
gal_type_bit_string 241
gal_type_from_name 241
gal_type_from_string 242
gal_type_is_list 241
gal_type_max 241

gal_type_min 241
gal_type_name 240
gal_type_out 241
gal_type_sizeof 240
gal_type_string_to_number............ 242
gal_type_to_string............ 242
gal_wcs_angular_distance_deg 278
gal_wcs_copy 277
gal_wcs_decompose_pc_cdelt 278
gal_wcs_img_to_world 279
gal_wcs_on_tile 277
gal_wcs_pixel_area_arcsec2 279
gal_wcs_pixel_scale_deg............ 279
gal_wcs_read 277
gal_wcs_read_fitsptr 277
gal_wcs_warp_matrix 278
gal_wcs_world_to_img 279
GAL_ARITHMETIC_FLAGS_ALL............ 284
GAL_ARITHMETIC_FREE 284
GAL_ARITHMETIC_INPLACE............ 284
GAL_ARITHMETIC_NUMOK 284
GAL_ARITHMETIC_OP_ABS 286
GAL_ARITHMETIC_OP_AND 285
GAL_ARITHMETIC_OP_BITAND............ 287
GAL_ARITHMETIC_OP_BITLSH............ 287
GAL_ARITHMETIC_OP_BITNOT............ 287
GAL_ARITHMETIC_OP_BITOR............ 287
GAL_ARITHMETIC_OP_BITRSH............ 287
GAL_ARITHMETIC_OP_BITXOR............ 287
GAL_ARITHMETIC_OP_DIVIDE............ 285
GAL_ARITHMETIC_OP_EQ 285
GAL_ARITHMETIC_OP_GE 285
GAL_ARITHMETIC_OP_GT 285
GAL_ARITHMETIC_OP_ISBLANK............ 285
GAL_ARITHMETIC_OP_LE 285
GAL_ARITHMETIC_OP_LOG 286
GAL_ARITHMETIC_OP_LOG10............ 286
GAL_ARITHMETIC_OP_LT 285
GAL_ARITHMETIC_OP_MAX 286
GAL_ARITHMETIC_OP_MAXVAL............ 286
GAL_ARITHMETIC_OP_MEAN............ 286
GAL_ARITHMETIC_OP_MEANVAL............ 286
GAL_ARITHMETIC_OP_MEDIAN............ 286
GAL_ARITHMETIC_OP_MEDIANVAL 286
GAL_ARITHMETIC_OP_MIN 286
GAL_ARITHMETIC_OP_MINUS............ 285
GAL_ARITHMETIC_OP_MINVAL............ 286
GAL_ARITHMETIC_OP_MODULO............ 287
GAL_ARITHMETIC_OP_MULTIPLY 285
GAL_ARITHMETIC_OP_NE 285
GAL_ARITHMETIC_OP_NOT 285
GAL_ARITHMETIC_OP_NUM 286
GAL_ARITHMETIC_OP_NUMVAL............ 286
GAL_ARITHMETIC_OP_OR 285
GAL_ARITHMETIC_OP_PLUS 285
GAL_ARITHMETIC_OP_POW 287
GAL_ARITHMETIC_OP_SQRT 286
GAL_ARITHMETIC_OP_STD 286
GAL_ARITHMETIC_OP_STDVAL............ 286
GAL_ARITHMETIC_OP_SUM 286
GAL_ARITHMETIC_OP_SUMVAL............ 286
GAL_ARITHMETIC_OP_TO_FLOAT32 287
GAL_ARITHMETIC_OP_TO_FLOAT64 287
GAL_ARITHMETIC_OP_TO_INT16 287
GAL_ARITHMETIC_OP_TO_INT32 287
GAL_ARITHMETIC_OP_TO_INT64 287
GAL_ARITHMETIC_OP_TO_INT8............ 287
GAL_ARITHMETIC_OP_TO_UINT16 287
GAL_ARITHMETIC_OP_TO_UINT32 287
GAL_ARITHMETIC_OP_TO_UINT64 287
GAL_ARITHMETIC_OP_TO_UINT8 287
GAL_ARITHMETIC_OP_WHERE............ 286
GAL_BINARY_TMP_VALUE 307
GAL_BLANK_FLOAT32............ 244
GAL_BLANK_FLOAT64............ 244
GAL_BLANK_INT16 243
GAL_BLANK_INT32 243
GAL_BLANK_INT64 244
GAL_BLANK_INT8 243
GAL_BLANK_SIZE_T............ 244
GAL_BLANK_STRING............ 244
GAL_BLANK_UINT16............ 243
GAL_BLANK_UINT32............ 243
GAL_BLANK_UINT64............ 244
GAL_BLANK_UINT8 243
GAL_CONFIG_BIN_OP_FLOAT32............ 235
GAL_CONFIG_BIN_OP_FLOAT64............ 235
GAL_CONFIG_BIN_OP_INT16............ 235
GAL_CONFIG_BIN_OP_INT32............ 235
GAL_CONFIG_BIN_OP_INT64............ 235
GAL_CONFIG_BIN_OP_INT8............ 235
GAL_CONFIG_BIN_OP_UINT16............ 235
GAL_CONFIG_BIN_OP_UINT32............ 235
GAL_CONFIG_BIN_OP_UINT64............ 235
GAL_CONFIG_BIN_OP_UINT8............ 235
GAL_CONFIG_HAVE_LIBGIT2............ 234
GAL_CONFIG_HAVE_PTHREAD_BARRIER............ 234
GAL_CONFIG_HAVE_WCSLIB_VERSION............ 234
GAL_CONFIG_SIZEOF_SIZE_T............ 235
GAL_CONFIG_VERSION............ 234
GAL_DIMENSION_FLT_TO_INT............ 254
GAL_DIMENSION_NEIGHBOR_OP............ 255
GAL_FITS_MAX_NDIM............ 268
GAL_POLYGON_MAX_CORNERS............ 299
GAL_POLYGON_ROUND_ERR 299
GAL_STATISTICS_BINS_INVALID 302
GAL_STATISTICS_BINS_IRREGULAR............ 302
GAL_STATISTICS_BINS_REGULAR 302
GAL_STATISTICS_MODE_GOOD_SYM 302
GAL_STATISTICS_SIG_CLIP_MAX_CONVERGE 302
GAL_STATISTICS_SORTED_DECREASING 302
GAL_STATISTICS_SORTED_INCREASING 302
GAL_STATISTICS_SORTED_NOT............ 302
GAL_TABLE_DEF_PRECISION_DBL 281
GAL_TABLE_DEF_PRECISION_FLT 281
GAL_TABLE_DEF_PRECISION_INT 281

GAL_TABLE_DEF_WIDTH_DBL 281
GAL_TABLE_DEF_WIDTH_FLT 281
GAL_TABLE_DEF_WIDTH_INT 281
GAL_TABLE_DEF_WIDTH_LINT 281
GAL_TABLE_DEF_WIDTH_STR 281
GAL_TABLE_DISPLAY_FMT_DECIMAL 281
GAL_TABLE_DISPLAY_FMT_EXP 281
GAL_TABLE_DISPLAY_FMT_FLOAT 281
GAL_TABLE_DISPLAY_FMT_GENERAL 281
GAL_TABLE_DISPLAY_FMT_HEX 281
GAL_TABLE_DISPLAY_FMT_OCTAL 281
GAL_TABLE_DISPLAY_FMT_STRING 281
GAL_TABLE_DISPLAY_FMT_UDECIMAL 281
GAL_TABLE_FORMAT_AFITS 282
GAL_TABLE_FORMAT_BFITS 282
GAL_TABLE_FORMAT_INVALID 282
GAL_TABLE_FORMAT_TXT 282
GAL_TABLE_SEARCH_COMMENT 282
GAL_TABLE_SEARCH_INVALID 282
GAL_TABLE_SEARCH_NAME 282
GAL_TABLE_SEARCH_UNIT 282
GAL_TILE_PARSE_OPERATE 292
GAL_TXT_LINESTAT_BLANK 279
GAL_TXT_LINESTAT_COMMENT 279
GAL_TXT_LINESTAT_DATAROW 279
GAL_TXT_LINESTAT_INVALID 279
GAL_TYPE_BIT 239
GAL_TYPE_COMPLEX32 240
GAL_TYPE_COMPLEX64 240
GAL_TYPE_FLOAT32 240
GAL_TYPE_FLOAT64 240
GAL_TYPE_INT16 239
GAL_TYPE_INT32 240
GAL_TYPE_INT64 240
GAL_TYPE_INT8 239
GAL_TYPE_INVALID 239
GAL_TYPE_SIZE_T 240
GAL_TYPE_STRING 240
GAL_TYPE_STRLL 240
GAL_TYPE_UINT16 239
GAL_TYPE_UINT32 239
GAL_TYPE_UINT64 240
GAL_TYPE_UINT8 239

P

pthread_barrier_destroy 236
pthread_barrier_init 236
pthread_barrier_t 236
pthread_barrier_wait 236
pthread_barrierattr_t 236

Index

$

$HOME 61
$HOME/.local/etc/ 61

–

-- 55
--cite 56
--config=STR 57
--disable-guide-message 39
--disable-progname 37
--dontdelete 53
--enable-bin-op-alltypes 38
--enable-bin-op-float32 37
--enable-bin-op-float64 37
--enable-bin-op-int16 37
--enable-bin-op-int32 37
--enable-bin-op-int64 37
--enable-bin-op-int8 37
--enable-bin-op-uint16 37
--enable-bin-op-uint32 37
--enable-bin-op-uint64 37
--enable-bin-op-uint8 37
--enable-gnulibcheck 38, 46
--enable-guide-message=no 39
--enable-progname 37
--enable-progname=no 37
--enable-reentrant 27
--hdu=STR/INT 52
--help 48, 56, 75
--help output customization 75
--ignorecase 52
--keepinputdir 53, 77
--lastconfig 57
--log 59
--numthreads 59, 62
--numthreads=INT 59
--onlyversion=STR 58
--output 59
--output=STR 53
--prefix 39
--printparams 50, 56
--program-prefix 43
--program-suffix 43
--program-transform-name 43
--quiet 56
--searchin=STR 52
--setdirconf 57, 61
--setusrconf 57, 61
--tableformat=STR 53
--type=STR 53
--usage 48, 55, 74
--version 56
--without-pgplot 28
-? 56
-D 53
-h STR/INT 52
-I 52
-K 53
-mecube (ds9) 343
-N INT 59
-o STR 53
-P 56
-q 56
-s STR 52
-S 57
-t STR 53
-T STR 53
-U 57
-V 56

.

./.gnuastro/ 61
./configure 36, 40
./configure options 37
.bashrc 75, 213
.desktop 343

3

32-bit 235

6

64-bit 235

A

A4 paper size 45
A4 print book 45
ACS 139
Additions to Gnuastro 11
Adjacency matrix 309
Adobe systems 88
ADU 201, 212
Advanced camera for surveys 139
Advanced Camera for Surveys 140
Affine Transformation 141
Albert. A. Michelson 4
Amplifier 73
Announcements 11
Anonymous bug submission 10
Anscombe F. J. 2
Anscombe's quartet 2
ANSI C 320
Aperture blurring 149
Aperture photometry 176, 186
Argp argument parser 75, 328

ARGP_HELP_FMT 75
args.h 328
Arguments to programs 48
Array 255
ASCII plot 156
ASCII table, FITS 67
ASCII85 encoding 92
astprogname 5
Astronomical data format 87
Astronomical data suffixes 49
Astronomical Magnitude system 201
Asynchronous thread allocation 15, 102
Atmosphere 117
Atmosphere blurring 149
authors-cite.h 329
Auto-complete in the shell 43
Automatic configuration file writing 60
Automatic output file names 77
Automatically created build files 33
Available number of threads 62
Average 303
Average, weighted 117
AWK 94, 96, 117, 157, 283

B

Background flux 151, 211
Background flux gradients 211
Background pixels 169
Backup 44
Best use of CPU threads 62
Bi-linear interpolation 142
Bias current 73
Bicubic interpolation 142
Bin width, histogram 148
Binary datasets 307
Binary image 88, 169
Binary table, FITS 68
Bit 64
bit-32 235
bit-64 235
Bitwise Or 284
Black and white image 88
blank color channel 88
Blank data 248
Blank pixel 100, 112
Blur image 117, 196
Blurring 149
Book formats 74
Bootstrapping 33
Border on an image 91
Breadth first search 196, 308
Brightness 198, 201
Buffers (Emacs) 323
Bug 9, 333
Bug reporting 9
Bug tracker 10
bug-gnuastro@gnu.org 10
Build 1
Build individual profiles 205
Build tree 332
Building from source 26
Byte 64

C

C Pre-Processor 232
C programming language 320
C++ programming language 320
C, plotting 344
C: restrict 246
C99 250
Cache, system 63
Camera 142
CANDELS 178
CCD 73, 139
Central management 333
CFITSIO 26, 268
CFITSIO version on outputs 78
Change converted pixel values 92
Channel 73
Charge-coupled device 139
Check 1
Check center of crop 106
Checking detection algorithms 195
Checking tests 45
Citation information 329
CLI: command-line user interface 7
CLI: repeating operations 8
Clump magnitude limit 177
CMYK 89
Colorspace 89
Colorspace, grayscale 89
Colorspace, transformation 90
Command-line arguments 48
Command-line help 74
Command-line options 48
Command-line scroll 75
Command-line searching text 75
Command-line token separation 48
Command-line user interface 7
Command-line, long outputs 75
Command-line, viewing full book 76
Comments 23
Commutative property 142
Comoving distance 219
Compare Moffat and Gaussian 197
Compare Poisson and Gaussian 211
Compile 1
Compiled PostScript 88
Compiling from source 26
Completeness 179
Complex numbers 137
Compression quality in JPEG 92
Configuration file directories 60
Configuration file format 59

Configuration file precedence 60
Configuration file suffix 59
Configuration files 51, 59
Configuration files, system wide 61
Configuration files, writing 60
Configuration, not finding library 46
Configure options 36
Configure options particular to Gnuastro 37
Configuring 36
Connected components 308
Connectivity 307
Convenient book formats 74
Convention for program source 326
Converting data formats 87
Converting image formats 87
ConvertType (`astconvertt`) 87
Convex Hull 299
Convolution 117, 119, 196
Convolution kernel 274
Cookbook 14
Coordinate transformation 139
Coordinates, homogeneous 140
Copyright 4
Correlated noise 181
Correlation 119
Cosmic ray removal 151
Cosmic rays 117, 139, 149, 151, 154
COSMOS survey 97
Counting error 210
Counting from zero 51
Counts 201, 212
`CPPFLAGS` 46, 226
CPU threads 62, 207
CPU threads, number 59
CPU threads, set number 59
CPU, using all threads 62
Crop (`astcrop`) 97
Crop a given section of image 100
Crop part of image 97
Crop section format 100
Cumulative Frequency Plot 148
Customize `--help` output 75
Customize executable names 43
Customizing installation 36

D

Data 153
Data format conversion 87
Data structures 225
Data type 246
Data's depth 178
Dataset: binary 307
de Vaucouleur profile 198
Debug 233
Debugging 332
Default executable search directory 40
Default library search directory 41
Default option values 51, 59
Define section to crop 100
Dependencies, Gnuastro 26
Depth 177
Detached threads 238
Detection 117, 163
Detections false 179
Detector 142
Development packages 46
Diffraction limited 196
Dilation 308
Directory, install 41
Discrete Fourier transform 137
Distortion, optical 139
Distribution mode 152
Douglas Rushkoff 2
Drizzle 143
Dynamic linking 227

E

Edges, image 142
Edwin Hubble 14
Effective radius 198
Efficient use of CPU threads 62
Ellipse 195
Elliptical distance 196
Elliptical galaxies 16
Emacs buffers 323
Encapsulated PostScript 88
Environment 40
Environment variable, `HOME` 40
Environment variables 39, 40, 213
`eog` 16
EPS 88
Erosion 164, 170, 308
Error, floating point round-off 137
`etc` 59
Exact area resampling 143
Executable names 43
Eye of GNOME 16

F

False detections 179
Feature request 333
Feature requests 11
File I/O 44
File operations 80
File system Hierarchy Standard 59
file systems, tmpfs 44
first-in-first-out 256, 271
FITS 268
FITS image viewer 343
FITS standard 26, 144, 246
FITS Tables 67
Fitting 195
Flip coordinates 140

Floating point error . . . 182
Floating point round-off error . . . 137
`FLT` . . . 50
Flux . . . 201
Flux to magnitude conversion . . . 201
Foreground pixels . . . 169
Fortran . . . 246
Fourier spectrum . . . 137
Free software . . . 4
Free Software Foundation . . . 335
FSF . . . 335
Full Width at Half Maximum . . . 197
Function gradient over pixel area . . . 199
Function groups . . . 325
Functions for user interface . . . 328
FWHM . . . 197

G

Gain . . . 201, 212
Galaxy profiles . . . 198
Galileo, G. . . . 3
Gaussian . . . 165, 186
Gaussian distribution . . . 153, 197
Gaussian FWHM . . . 197
GCC . . . 230
Gedit . . . 17, 23
General file operations . . . 80
Generalized de Vaucouleur profile . . . 198
Gérard de Vaucouleurs . . . 198
Git . . . 28, 32, 283, 311
GNOME 2 . . . 344
GNOME 3 . . . 7, 343
GNU Astronomy Utilities (Gnuastro) . . . 1
GNU Autoconf . . . 30, 34, 35, 36, 331
GNU Autoconf Archive . . . 30, 33
GNU Automake . . . 30, 34, 331
GNU AWK . . . 94, 95, 96, 117, 157, 283
GNU Bash . . . 8, 16, 41, 85, 322
GNU Binutils . . . 227
GNU build system . . . 26, 34, 41, 44, 226, 331
GNU C library . . . 7, 29, 34, 38, 44, 76, 185, 228, 324, 328
GNU coding standards . . . 1, 323, 325
GNU Compiler Collection . . . 7, 230, 232, 323
GNU Coreutils . . . 62, 322
GNU CPP . . . 232
GNU Emacs . . . 8, 16, 17, 23, 45, 76, 325, 326
GNU Free Documentation License . . . 5, 346
GNU General Public License . . . 5
GNU Grep . . . 75, 84, 221
GNU help2man . . . 30
GNU Info . . . 76
GNU Libtool . . . 30, 34, 47, 227, 229, 231, 232, 331
GNU Make . . . 64, 231
GNU Parallel . . . 16, 63
GNU Portability Library (Gnulib) . . 29, 33, 38, 46, 324
GNU Savannah . . . 333
GNU Scientific Library . . . 26, 213, 301
GNU software documentation . . . 76
GNU style options . . . 50
GNU Tar . . . 1
GNU Texinfo . . . 4, 34, 45, 47
GNU/Linux . . . 7
Gnuastro coding convention . . . 323
Gnuastro common options . . . 52
Gnuastro major version number . . . 6
Gnuastro program structure convention . . . 326
Gnuastro project page . . . 10
Gnuastro test scripts . . . 332
Gnulib: GNU Portability Library . . . 29, 33, 38, 46, 324
GPL Ghostscript . . . 28, 29, 47, 88
Gradient over pixel area . . . 199
Gradients in background flux . . . 211
Graphic user interface . . . 7
Gravitational lensing . . . 139
Grayscale . . . 89
Groups of similar functions . . . 325
GUI: graphic user interface . . . 7
GUI: repeating operations . . . 8
Gzip . . . 32

H

Halted program . . . 9
HDD . . . 44
HDU . . . 49, 52, 82
Header data unit . . . 49, 52
Header file . . . 323
Help . . . 74
help-gnuastro mailing list . . . 77
`help-gnuastro@gnu.org` . . . 77
Hexadecimal encoding . . . 92
Histogram . . . 148, 152
`HOME` . . . 40
`HOME/.local/` . . . 40
Homogeneous coordinates . . . 140
Homography . . . 141
Hubble Space Telescope . . . 73, 97, 139, 140
Hyper Suprime-Cam . . . 73

I

Image . . . 89
Image blurring . . . 196
Image edges . . . 142
Image format conversion . . . 87
Image mosaic . . . 97, 139
Image noise . . . 210
Image tiles . . . 97
Image transformations . . . 195
ImageMagick . . . 31
Imaging surveys . . . 97
Immediate neighbors . . . 307

Inconsistent results . . . 9
Individual profiles . . . 205
info-gnuastro@gnu.org . . . 35
INFOPATH . . . 41
Input/Output, file . . . 44
Inside-out construction . . . 196, 199
Install directory . . . 41
Install with no super-user access . . . 39
Installation . . . 25
Installation, customizing . . . 36
Installed help methods . . . 74
Instrumental noise . . . 212
Integer, Signed . . . 64
Integration over pixel . . . 199
Integration to infinity . . . 202
Internal default value . . . 59
Internally stored option value . . . 62
Interpolation . . . 142
Interpolation, bi-linear . . . 142
Interpolation, bicubic . . . 142
Intervals, histogram . . . 148
INT . . . 50
ISO C90 . . . 320
Issue . . . 333

J

Jaynes E. T. . . . 4
JPEG compression quality . . . 92
JPEG format . . . 29, 87

K

Ken Thomson . . . 4
Kernel, convolution . . . 117, 274
Kernighan, Brian . . . 320

L

Labeling . . . 163
Large astronomical images . . . 97
last-in-first-out . . . 256, 271
LaTeX . . . 30, 88, 322
Lawrence Livermore National Laboratory . . . 235
LD_LIBRARY_PATH . . . 41, 46, 345
LDFLAGS . . . 46
Learning GNU Info . . . 76
Lensing simulations . . . 195
less . . . 75
libgit2 . . . 28, 311
libjpeg . . . 29
Library search directory . . . 41
Library: shared . . . 227
Limit, object/clump magnitude . . . 177
Linear spatial filtering . . . 118
Linked list . . . 255, 271
Linking: dynamic . . . 227
Linking: Static . . . 227
Linux . . . 7
Linux kernel . . . 44
Long option abbreviation . . . 50
Long outputs . . . 75
Lord Kelvin . . . 4
Low level programming . . . 321
Luminosity . . . 201
Lzip . . . 32

M

Macro . . . 225
Magnitude zero-point . . . 201
Magnitude, object/clump detection limit . . . 177
Magnitude, upper limit . . . 181
Magnitudes from flux . . . 201
Mailing list archives . . . 10, 77
Mailing list: bug-gnuastro . . . 10, 333
Mailing list: gnuastro-commits . . . 335, 337, 338
Mailing list: gnuastro-devel . . . 334, 335
Mailing list: help-gnuastro . . . 77
Mailing list: info-gnuastro . . . 5, 11, 35
main function . . . 327
Main parameters C structure . . . 327
main.c . . . 327
main.h . . . 327
Major version number . . . 5
Make . . . 64
make check . . . 45
MakeProfiles (astmkprof) . . . 195
Making a distribution package . . . 333
Making profiles pixel by pixel . . . 196
Man pages . . . 76
Management hub . . . 333
Mandatory arguments . . . 48, 74
MANPATH . . . 41
Mathematical morphology . . . 308
Matplotlib, Python . . . 322, 344
Matrix . . . 140
Matrix multiplication . . . 142
Matrix, adjacency . . . 309
Maximum . . . 303
Mean . . . 303
Median . . . 153, 303
Meta-data . . . 80
Metacharacters . . . 48
Metacharacters on the command-line . . . 49
Michelson, Albert. A. . . . 4
Minimum . . . 303
Minor version number . . . 5
Mixing pixel values . . . 117, 142
Möbius, August. F. . . . 140
mock.fits . . . 45
Mode . . . 153
Mode of a distribution . . . 152
Modeling . . . 195
Modeling stars . . . 198
Modifying print book . . . 45

Modularity . . . 223
Moffat beta . . . 197
Moffat function . . . 197
Moffat FWHM . . . 197
Moments . . . 181
Monte carlo integration . . . 199
Mosaicing . . . 97, 139
Multi-threaded programs . . . 62
Multiextention FITS . . . 343
Multiple file opening, reentrancy . . . 27
Multiplication, matrix . . . 142
Multiplication, Matrix . . . 140
Multithreaded programming . . . 235

N

Names of executables . . . 43
Names, customize . . . 43
Names, programs . . . 5
NaN . . . 70, 110, 146, 162, 169, 188, 243, 244, 274
Navigating source files . . . 326
Necessary parameters . . . 59
Neighborhood . . . 117
Neighbors, immediate . . . 307
No access to super-user install . . . 39
Noise . . . 153, 210
Noise simulation . . . 211
Noise, instrumental . . . 212
Non-commutative operations . . . 142
Normalizing histogram . . . 148
`nproc` . . . 62
Number . . . 303
Number of CPU threads to use . . . 59
Number of threads available . . . 62
Number, version . . . 5
Numbers, complex . . . 137
Numbers, psudo-random . . . 212
Numbers, random . . . 212

O

Object magnitude limit . . . 177
Object oriented programming . . . 320
On/Off options . . . 50
Online help . . . 74
Opening (Mathematical morphology) . . . 308
Opening multiextention FITS . . . 343
OpenMP . . . 235
Operations on files . . . 80
Operations, non-commutative . . . 142
Operator, structure de-reference . . . 327
Optical distortion . . . 139
Optimization . . . 232, 332
Optimization flag . . . 323
Option values . . . 50
Optional and mandatory tokens . . . 74
Options . . . 117
Options common to all programs . . . 52
Options to programs . . . 48
Options, abbreviation . . . 50
Options, GNU style . . . 50
Options, on/off . . . 50
Options, repeated . . . 51
Options, short (`-`) and long (`--`) . . . 50
Order in search directory . . . 42
Output file names, automatic . . . 77
Output FITS headers . . . 78
Output, wrong . . . 9
Oversample . . . 20
Oversampling . . . 200

P

`p` . . . 327
Package managers . . . 26
Paper size, A4 . . . 45
Paper size, US letter . . . 45
Parametric PSFs . . . 197
`PATH` . . . 40
PDF . . . 88
permutation . . . 301
PGFplots in TeX or LaTeX . . . 322, 344
PGPLOT . . . 344
Phase angle . . . 137
photo-electrons . . . 152
Photoelectrons . . . 142
Photometry, aperture . . . 176, 186
Photon counting noise . . . 210
Picture element . . . 142
Pipe . . . 75
Pixel . . . 142
Pixel by pixel making of profiles . . . 196
Pixel mixing . . . 117, 142, 143
Pixelated graphics . . . 87
Pixels . . . 89
Plain text . . . 89
Plotting directly in C . . . 344
PNG standard . . . 90
Point pixels . . . 142
Point source . . . 196
Point Spread Function . . . 196
Poisson distribution . . . 211
Portable Document format . . . 88
Position angle . . . 183
POSIX Threads . . . 237
POSIX threads . . . 236
POSIX threads library . . . 235
Postage stamp images . . . 97
Postfix notation . . . 108
PostScript . . . 88
PostScript vs. PDF . . . 88
Pre-Processor . . . 224
Pre-processor macros . . . 225
Precedence, configuration files . . . 60
`prefix/etc/` . . . 61
Primary colors . . . 89

printf ... 281
Probability density function ... 148, 152
Profiles, galaxies ... 198
progname.c, progname.h ... 329
prognameparams ... 327
Program crashing ... 9
Program names ... 5
Program structure convention ... 326
Programming, low level ... 321
ProgramName ... 5
Projective transformation ... 141
Proper distance ... 219
PSF ... 196
PSF image size ... 196
PSF over-sample ... 200
PSF width ... 197
PSF, Moffat compared Gaussian ... 197
Psudo-random numbers ... 212
pthread ... 62
pthread_barrier ... 236
Public domain ... 4
Purity ... 179
Puzzle solving scientist ... 4
Python Matplotlib ... 322, 344
Python programming language ... 320

Q

Quality of compression in JPEG ... 92
Quantile ... 156, 169, 304

R

Radial profile on ellipse ... 196
Radius, effective ... 198
Random numbers ... 212
Raster graphics ... 87
Readout noise ... 212
Redirection of output ... 75
Reentrancy, multiple file opening ... 27
Remembering options ... 74
Remote operation ... 9
Removing ast from executables ... 44
Repeated options ... 51
Report a bug ... 333
Reproducibility ... 189
Reproducible bug reports ... 10
Reproducible results ... 8
Resampling ... 142
Resource heavy operations ... 9
restrict ... 246
Results, wrong ... 9
Reverse Polish Notation ... 108
RGB ... 89
Ritchie, Dennis ... 320
Root access, not possible ... 39
Root parameter structure ... 327
Rotation of coordinates ... 140
Round-off error ... 137, 299

S

Sampling ... 142, 199
SAO ds9 ... 104, 343
Save output to file ... 75
Saving binary image ... 88
Scaling ... 139
Scientist, puzzle solver ... 4
Scripts, startup ... 41
Scroll command-line ... 75
Search directory for executables ... 40
Search directory order ... 42
Searching text ... 75
Second moment ... 182
Section of an image ... 97
Secure shell ... 9
SED, stream editor ... 44
Seed, psudo-random numbers ... 212
Segmentation ... 163
Separating tokens on the command-line ... 48
Sérsic index ... 198
Sérsic profile ... 198
Sérsic, J. L. ... 198
Setting output file names automatically ... 77
Setting PATH ... 40
Shared library ... 227
Shared library versioning ... 228
Shear ... 140
Shell ... 7
Shell auto-complete ... 43
Shell script ... 6
Shell variables ... 39
Shift + PageUP and Shift + PageDown ... 75
sigma-clipping ... 153
Signal ... 153
Signal to noise ratio ... 139, 142
Signed integer ... 64
Simulating noise ... 211
Simultaneous multithreading ... 62
Single channel CMYK ... 90
size_t ... 265, 266
Skewed Poisson distribution ... 211
Sky ... 150
Sky value ... 151, 152, 211
Software bug ... 9
Source code building ... 26
Source code compilation ... 26
Source file navigation ... 326
Source tree ... 332
Source, uncompress ... 1
Spectrum, Fourier ... 137
Spiral galaxies ... 16
Spread of a point source ... 196
SSD ... 44
SSH ... 9

Standard deviation . . . 182, 303
Standard, FITS . . . 246
Stars, modeling . . . 198
Startup scripts . . . 41, 213
Static document description format . . . 88
Static linking . . . 227
Statistical analysis . . . 2
Stitch multiple images . . . 97
Stream editor, SED . . . 44
Stroustrup, Bjarne . . . 4, 320
Structure de-reference operator . . . 327
Structures . . . 225
`STR` . . . 50
Subaru Telescope . . . 73
Submit new tracker item . . . 10
Suffixes, astronomical data . . . 49
Suffixes, EPS format . . . 88
Suffixes, JPEG images . . . 87
Suffixes, PDF format . . . 88
Suffixes, plain text . . . 89
Sum . . . 303
Sum for total flux . . . 201
Superuser, not possible . . . 39
Support request manager . . . 10
Surface brightness . . . 178
Symbolic link . . . 43
System Cache . . . 63
System wide configuration files . . . 61

T

Tables FITS . . . 67
Tabs are evil . . . 325
Task tracker . . . 10
Test . . . 1
Test scripts . . . 332
Tests, only one passes . . . 46
Tests, running . . . 45
`tests/during-dev.sh` . . . 331
TeX . . . 47, 88
TeX Live . . . 30
Threads, CPU . . . 207
Thresholding . . . 307
Tilde expansion as option values . . . 51
tmpfs file system . . . 44
`tmpfs-config-make` . . . 331, 332
Token separation . . . 48
Top processing source file . . . 329
Top root structure . . . 327
Tracker . . . 10, 333
Trailing space . . . 325
Transform image . . . 195
Transformation, affine . . . 141
Transformation, projective . . . 141
Truncation radius . . . 201
Tutorial . . . 14
Type . . . 64

U

`ui.c` . . . 328
`ui.h` . . . 328
Uncompress source . . . 1
Undetected objects . . . 211
Unsigned integer . . . 64
Upper limit magnitude . . . 181
US letter paper size . . . 45
Usage pattern . . . 74
User interface functions . . . 328
Using CPU threads . . . 62
Using multiple CPU cores . . . 62
Using multiple threads . . . 62

V

Values to options . . . 50
Variance . . . 182
Variation of background flux . . . 211
Vector graphics . . . 88
Version control . . . 9, 32
Version control systems . . . 28
Version number . . . 5
Versioning: Shared library . . . 228
Viewing trackers . . . 10
Virtual console . . . 8
`void *` . . . 246

W

Wall-clock time . . . 62
WCS . . . 28
WCSLIB . . . 28, 147
Weighted average . . . 117
WFC3 . . . 139
White space character . . . 59
Wide Field Camera 3 . . . 139, 140
William Thomson . . . 4
World Coordinate System . . . 28, 147
Writing configuration files . . . 60
Wrong output . . . 9
Wrong results . . . 9

X

XDF . . . 178

Z

Zero-point magnitude . . . 201