# Making Canonical Workflow Building Blocks interoperable across workflow languages

Stian Soiland-Reyes, Genís Bayarri, Pau Andrio, Robin Long, Douglas Lowe, Ania Niewielska, Adam Hospital, Paul Groth

#### Abstract

We introduce the concept of *Canonical Workflow Building Blocks* (**CWBB**), a methodology of describing and wrapping computational tools, in order for them to be utilized in a reproducible manner from multiple workflow languages and execution platforms. The concept is implemented and demonstrated with the BioExcel Building Blocks library (BioBB), a collection of tool wrappers in the field of computational biomolecular simulation. Interoperability across different workflow languages is showcased through a protein Molecular Dynamics setup transversal workflow, built using this library and run with 5 different Workflow Manager Systems (WfMS). We argue such practice is a necessary requirement for FAIR Computational Workflows and an element of Canonical Workflow Frameworks for Research (CWFR) in order to improve widespread adoption and reuse of computational methods across workflow language barriers.

## **1. Introduction**

The need for *reproducibility* of research software usage is well established [Stodden 2016, Leipzig 2020, Katz 2021], and adaptation of *workflow management systems* (WfMS) together with *software packaging and containers* [Möller 2017] have been proposed as key ingredients for making research software usage *FAIR* and reproducible [Cohen-Boulakia 2017, Grüning 2018a, Lamprecht 2020]. Recently it is also argued that computational workflows should also be treated as FAIR Digital Objects [De Smedt 2020] in their own right, with identifier, metadata and interoperability requirements [Goble 2020].

BioExcel, a European Centre of Excellence for Computational Biomolecular Research, has a particular focus on the boundary of molecular dynamics simulations and bioinformatics analytics using *High Performance Computing* (HPC) to approach Exascale performance, while also improving usability. The *BioExcel Building Blocks* (BioBB) [Hospital 2020] have been created as portable wrappers of open-source computational tools identified as useful for BioExcel workflows, forming several *families* of documented and interoperable operations that can be called from multiple workflow systems. This interoperability is shown with the BioBB demonstrator workflows, along with multiple tutorials and notebooks.

We propose that these building blocks and their families can themselves be considered *composite Digital Objects*: collections of software packages and their source code, guides and tutorials, as well as workflow management system integrations and workflow examples. In addition, the building blocks, as wrappers of upstream open source tools, benefit from and refer to the tools' existing documentation, support forums, academic publications and wider development context.

Given BioBB as a starting point, we define a generalized methodology of Canonical Workflow Building Blocks through the definition of a set of requirements and recommendations for how to formalize and develop a family of compatible computational tools as Digital Objects. These building blocks let researchers instantiate a Canonical Workflow in multiple workflow management systems, while also benefiting from the FAIR aspects of the CWBB Digital Objects.

#### 2. Methods

The <u>BioExcel Building Blocks</u> library [<u>Andrio 2019</u>], created and implemented within the <u>BioExcel CoE</u>, is a collection of portable wrappers of common biomolecular simulation tools. The library is designed to i) increase the *interoperability* between the tools wrapped; ii) *ease the implementation* of biomolecular simulation workflows; and iii) increase the *reusability and reproducibility* of the generated workflows. To achieve these main goals, the library was designed following the FAIR principles for research software development best practices [<u>Lamprecht 2020</u>]. The result is a collection of building block modules, divided in sets of tool wrappers focused on similar functionalities (e.g. Molecular Dynamics, Virtual Screening). Each of the modules is built from a combination of (i) software packaging (<u>Pip, BioConda</u>, BioContainers), (ii) documentation (<u>ReadTheDocs</u>), (iii) interactive tutorials (<u>Jupyter Notebooks</u>, <u>myBinder</u>), (iv) registry & findability (<u>bio.tools</u>, <u>BioSchemas</u>, <u>WorkflowHub</u>), (v) WfMS integration stubs (<u>CWL</u>, <u>Galaxy</u>, <u>PyCOMPSs</u>), (vi) source Code (<u>GitHub</u>) and (vii) REST APIs (<u>OpenAPI</u>, <u>Swagger</u>). Notably all building blocks follow the same pattern of installation, configuration and interaction.

Since the publication of the library, several <u>new building block modules</u> (Chemistry, Machine Learning, Amber, Virtual Screening, etc.) have been added, and the set of operations for the existing BioBB families have been expanded. While we previously provided curated adapters (Figure 1) for running BioBB in workflow systems using CWL and PyCOMPSs, along with Galaxy Toolshed bindings, we have now started auto-generating these bindings, along with command line wrappers and REST web service APIs, using annotations within BioBB's Python docstrings as source. These annotations include sufficient information for a WfMS to launch a particular building block: *input* and *output* parameters (including *mandatory/optional* flags), compatible *formats* (including EDAM ontology formats), *example* files (essential for testing purposes), default values and *dependencies*. This ensures human-readable documentation, FAIR metadata and programmatic accessibility can be generated consistently and comparably.

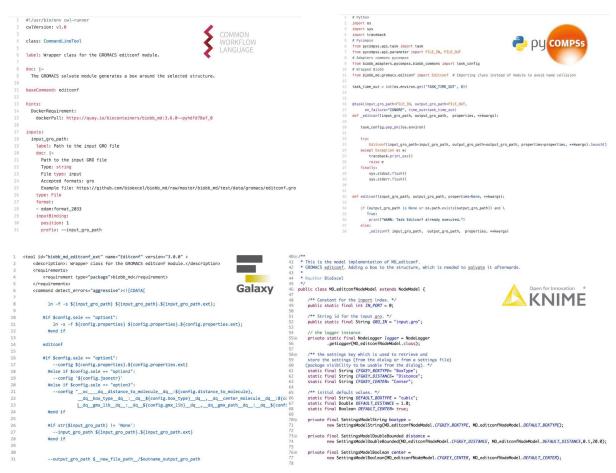


Figure 1: Code snippets for the BioBB WfMS bindings: CWL, PyCOMPSs, Galaxy and KNIME

The library is showcased through a collection of <u>demonstration workflows</u>. Jupyter Notebook has been highlighted as a valuable tool for reproducible scientific workflows [Beg 2021], which we recognize by primarily exposing the workflows as Jupyter Notebooks. This offers a graphical interactive interface, including documentation (integrated markdown) related to the workflow and the building blocks used, but also to the biomolecular simulation methods used in the pipeline. Moreover, as we have demonstrated with <u>Binder [Jupyter 2018]</u>, these workflows are reproducible across platforms, assisted by BioConda [Grüning 2018b] packaging of the building blocks and their software dependencies. The assembly of available demonstration workflows have been successfully used in the BioExcel CoE for dissemination and training events (e.g. <u>BioExcel Summer/Winter School</u>).

We can observe that workflow building blocks such as BioBB are necessarily composed of a comprehensive list of digital objects, encompassing source code, packaging, containerization, documentation, attributions, citations, registry entries, WfMS integrations and REST APIs.

Considering building blocks as *composite digital objects* in their own right - which along with their metadata, identifiers and operations can be called *canonical workflow building blocks* (**CWBB**) - allows for them to be considered as fundamental elements of FAIR Digital Objects for Computational Workflows that use the building blocks as operations, propagating information and resources that are essential for reproducibility, reuse and understanding by anyone discovering the workflow.

#### 2.1 Interoperability across different workflow languages

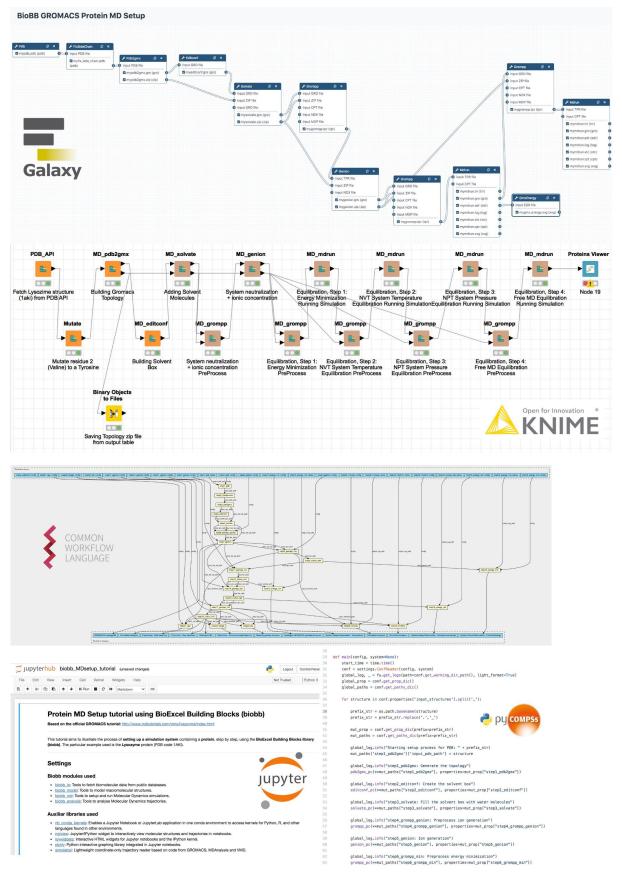
The concept of canonical workflow building blocks is showcased with the BioBB library using a transversal workflow present in many different computational biomolecular projects: a <u>Molecular Dynamics</u> (<u>MD</u>) protein setup. This workflow prepares a protein structure to be used as input for an MD simulation, going through a series of steps where the protein is completed (adding hydrogen and missing atoms), optionally introducing a residue mutation, then submerging the protein in a virtual box of water molecules with a particular ionic concentration, and finally energetically equilibrating the system (so that solvent and ions are well accommodated around the protein at the desired temperature).

This simulation process involves a non-negligible number of steps, using a variety of biomolecular tools. The BioBB library was used to assemble this workflow, interconnecting building blocks using Python functions (Jupyter Notebook, Command Line Interface), auto-generated bindings (Galaxy, CWL, PyCOMPSs) or manually generated bindings (KNIME). Corresponding workflows for the different WfMS can be found in WorkflowHub and graphical extracts can be seen in Fig. 2.

This example demonstrates how the same canonical building blocks can be used in different WfMS. Wrappers and tools executed behind the workflows are exactly the same, but the workflows are built using different WfMS, some of them in a graphical way (drag & drop, Galaxy, KNIME), some in a command line way (Jupyter Notebook, PyCOMPSs, CWL); workflows can be focused on short/interactive executions (Jupyter Notebook), or on High Throughput/High Performance Computing (HT-HPC) executions (PyCOMPSs); some of them prepared for a particular WfMS installation (Galaxy), others completely system-agnostic (CWL).

The current number of available bindings include Jupyter Notebook, PyCOMPSs, CWL, Galaxy and KNIME WfMS, apart from a recommended <u>command line</u> mechanism. Thanks to the extensive documentation added in the source code as Python docstrings, new bindings for available WfMS can be generated. However, it should be noted that this automatic generation of WfMS bindings is not always practically feasible. As an example, KNIME nodes require a complete Java skeleton code, as well as a definition of new data types for all inputs/outputs required, which makes their automatic generation a heavy and potentially error-prone task. Bindings for workflow languages with a *domain-specific language* (DSL) for tool definitions (e.g. Galaxy, CWL) can on the other hand be generated in a more straightforward fashion.

The transversal protein MD setup workflow was chosen as a real example that is readily understandable by domain experts. More <u>complex pipelines</u> involving a broader set of wrapped biomolecular tools have been developed using the BioBB library, primarily as Jupyter Notebooks. A selection of these will similarly be assembled for different WfMS using the auto-generated bindings and uploaded to the <u>WorkflowHub repository</u>.



**Figure 2**: Protein MD Setup transversal workflow, assembled in with 5 different workflow managers using BioBB canonical building blocks. From top-left: Galaxy [10.48546/workflowhub.workflow.194.1], KNIME

#### 3. Discussion

Early work on libraries of workflows fragments include Web Service-based approaches where tools are wrapped and exposed using common, interoperable data types in <u>BioMoby</u> [Biomoby 2008] for bioinformatics and similarly <u>caBIG</u> [Saltz 2006] for cancer genomics. While these efforts were interoperable across WfMSs they required a large up-front investment in agreeing to and adapting native data to common RDF or XML representations.

The notion of *abstract workflows* [Garijo 2011], structural workflow descriptions separated from their concrete execution realizations and augmented with Linked Data annotations, have been emphasized as essential for reuse and consistency across workflow systems. Identifying *common motifs* for workflow operations [Garijo 2014] (e.g. Data preparation, Format transformation, Filter, Combine) are important to simplify and understand otherwise fine-grained workflow provenance traces.

Most other efforts to standardize a set of disparate analytical tools have been done within the scope of a single WfMS, allowing customized user interaction, data visualization, configuration and findability, for instance <u>Taverna components</u> had prototypical building blocks [<u>Giovanni 2016</u>] which were instantiated at runtime by reference from a registry. <u>KNIME components and metanodes</u>, shared on the <u>KNIME Hub</u> are frequently designed to be interoperable, but with a perhaps weaker notion of component families. The <u>Galaxy</u> toolshed [<u>Blankenberg 2014</u>] is likewise populated with different sets of tool wrappers that are largely made to be interoperable within a category.

The Common Workflow Language (CWL) [Crusoe 2021] has a strong emphasis on interoperable command line tool descriptions, with <u>support for containers</u> and Conda packaging, as well as <u>support for FAIR metadata</u> like contributors, license and EDAM ontology type annotations. With multiple leading workflow engines now supporting CWL, and experimental Galaxy support, this seems perhaps the most promising candidate for both making and describing canonical workflow building blocks, however we've identified a few stumbling blocks.

One obvious challenge is that the implementing WfMS needs to have CWL support, along with support for either containers or Conda packaging to find the described executables. While it is possible to run a CWL tool directly using a #1/usr/bin/env cwl-runner shebang on POSIX systems, this still requires pre-installation and possibly configuration of a CWL engine like cwltool, which themselves have multiple dependencies and cannot easily be used as a container, because to execute the tool it would need *nested containers* which are not generally permitted for security reasons. Within the CWL community it was originally envisioned that a wider set of workflow systems would adopt CWL for tool description/execution, with a subset implementing full CWL workflow support. This would allow shared community effort for describing tools, say in the <u>Common</u> Workflow Library, rather than each WfMS needing to duplicate this tool wrapping in separate repositories and languages. However, with the exception of experimental tool support in Galaxy, in practice all CWL implementers have gone for full workflow support.

Another challenge is that making a set of building blocks frequently requires the use of *shims*, for instance file conversion, small search/replace operations or file renames. In a CWL approach these can either be performed with an <u>Expression</u> using JavaScript snippets which only has limited access to file content, or as an additional workflow step added before or after the main tool step. This could then be nested as a subworkflow, similar to KNIME's metanodes, and would also allow different containers or packages for any pre- or post-steps. This however becomes harder to access from a non-CWL WfMS because of lack of control over configuration/execution options for the now nested CWL tools, and in practice requires CWL-supporting WfMS to implement full CWL Workflow support.

It is worth mentioning that it would be possible to generate WfMS-specific tool bindings from CWL (e.g. as demonstrated with <u>cwl2script</u> for Bash, <u>gxargparse</u> for Galaxy, <u>cwl2wdl</u> for WDL), although this necessitates constraining the tool and workflow definitions to a limited subset of CWL.

For the main BioBB building blocks we implemented <u>demonstrator workflows</u> that highlight how the tools should be used in different workflow management systems; each having a primary exemplar using Jupyter Notebook, which can be explored interactively using the <u>BioExcel myBinder</u>. If we consider the abstract

demonstrator workflows as *canonical workflows* they are therefore very much active objects, but can also be seen as *workflow templates*, as any real use case will need to specialize the workflow to tweak parameters, data selection etc.

We therefore also now provide such workflow templates for multiple WfMS, including CWL, PyCOMPSs and Galaxy, which are fairly disparate workflow languages, yet by the use of the same canonical workflow building blocks (which again invoke the same software binaries), all effectively are instantiations of the same canonical workflow.

One challenge found is how to publish such canonical workflows in registries like the <u>WorkflowHub</u>. The hub supports the registration of Digital Objects in the form of RO-Crate [<u>Soiland-Reves 2021</u>], with abstract CWL for describing the canonical workflow template, along with direct references to the workflow's GitHub repository, however this results in one registration entry per workflow language, which are not otherwise related and may not even share the same source code repository. Thus we've identified the need for adding an overall *canonical workflow entry*, which can bring in workflow documentation and references shared across WfMS implementations, including a set of links to the more granular canonical workflow building blocks used by the workflow, but also to the individual WfMS implementations as separate digital objects.

A similar question of granularity applies at the workflow tool level [Möller 2017], particularly for Findability and Accessibility, as we can consider at lowest granularity the *scientific method* in general (e.g. any algorithm for sequence alignment), followed by an *application suite* (bio.tools entry, homepage, documentation), instantiated as a particular *software installation* (Debian package, Docker container) with its dependencies at same level, which includes one or more *software executables* (a particular binary, a running service service), providing at the highest granularity level the specific types of *software functionality* (a particular mode of operation, choice of analysis) for instance using different command line flags.

For canonical workflow building blocks, with a focus on pluggable composability, this is mainly defined at the lowest granularity level of software functionality, operations that can be added to a workflow. This is indeed the level WfMS tool definitions are done, e.g. a CWL Command Line Tool specifies a particular way to run a particular software binary. However to be a CWBB the building block needs to additionally rely on the lower granularity levels, particularly to support multiple options for interoperable installation and execution, as well as metadata at the most general level, such as documentation and scholarly citations.

While workflow management systems typically only operate at the highest granularity levels for execution details, and are frequently unaware of (or not exposing metadata at) the more general levels, we argue that in order for a Canonical Workflow [Wittenburg 2021] to follow and support FAIR principles for itself and its data, the workflow management system need to *propagate structured metadata* about the tools used by the workflow. We propose that in order to support the workflow's applicability to multiple WfMS, the tools themselves must also have a consistent packaging and formal description that enables consistent computational invocation.

At the most general level, a canonical workflow built using such CWBBs is even conceptually reproducible because the FAIR documentation of the workflow, through its canonical workflow building blocks, identifies how individual tools and software applications are composed, which in worst case can be rebuilt using different installation methods in a different WfMS, or in best case inspected to detect and cross-link the same canonical workflow appearing in different WfMS instantiations. This view of software as composition of other software typically also applies at individual tool level, which themselves depend on programming language runtimes, libraries, services and reference data.

## 4. Requirements for Canonical Workflow Building Blocks

Building on the experiences with BioBB, we here propose requirements and recommendations for establishing Canonical Workflow Building Blocks (CWBB) as implementations of *canonical steps* introduced for Canonical Workflow Frameworks for Research [Wittenburg 2021].

The core purpose of a CWBB is to wrap a command line tool or other software that can perform an operation as part of a computational workflow. As such, the general advice for making software workflow-ready applies [Brack 2021] (e.g. easy to install, documented, parallelisable, reproducible output), however a CWBB is

also permitted to make use of additional scripts or *shims* to further adapt a third-party tool for workflow use and for data interoperability across blocks.

The way tools are installed or invoked varies slightly across WfMS and operating systems, therefore a CWBB should provide multiple methods for distributing software; currently containers (Docker, Singularity) and distribution-independent packaging (e.g. Conda, Homebrew) are promising by having reproducible install recipes and a wide range of open source dependencies (e.g. Java, Python). Additionally building blocks should allow overriding execution paths, e.g. for use with HPC module system and hardware-optimized binaries.

The CWBBs should have sufficient annotations to be able to generate bindings for different WfMSs and REST APIs, e.g. parameter names and descriptions, types and default values; enumerators for options, file formats for inputs/outputs.

Building blocks should be grouped into families that are interoperable through common data structures and file formats, as well as having joint naming conventions for configuration options. A CWBB family should be released as a single version following <u>semantic versioning</u> rules, which should have a corresponding persistent identifier (PID).

Metadata for CWBBs should be captured following FAIR guidelines, and distributed as part of the block family and resolvable from the PID as a FAIR Digital Object. Metadata should include references to the CWBB software distributions (e.g. quay.io container URL) as well as attributions, citations and documentation for the wrapped tool.

Example workflows showing CWBB usage should be included in a WfMS-neutral language such as Jupyter Notebooks, which may have equivalent variants for each workflow binding. These workflows should be registered in a workflow registry like WorkflowHub or Dockstore, and assigned their own PIDs.

## 5. Conclusions

The proposed concept of Canonical Workflow Building Blocks can bridge the gap between FAIR Computational Workflows, interoperable reproducibility and for building canonical workflow descriptions to be used and described FAIRly across WfMSs.

The realization of CWBBs can be achieved in many ways, not necessarily using the Python programming language together with RO-Crate as explored here. In particular if the envisioned Canonical Workflow Frameworks for Research become established in multiple WfMSs with the use of FAIR Digital Objects, the different implementations will need to agree on object types, software packaging and metadata formats in order to reuse tools and provide interoperable reproducibility for canonical workflows. The Common Workflow Language shows promise as a way to gather execution details of tools together with their metadata and references, but needs further work to formalize practices to make CWL available as FAIR Digital Objects on their own or as part of an aggregate collection like RO-Crate.

## Acknowledgements

This work has been done as part of the BioExcel CoE (<u>https://www.bioexcel.eu/</u>), a project funded by the European Union contracts <u>H2020-INFRAEDI-02-2018 823830</u>, <u>H2020-EINFRA-2015-1 675728</u>. Additional work is funded through EOSC-Life (<u>https://www.eosc-life.eu/</u>) contract <u>H2020-INFRAEOSC-2018-2 824087</u>, and ELIXIR-CONVERGE (<u>https://elixir-europe.org/</u>) contract <u>H2020-INFRADEV-2019-2 871075</u>.

# **Author affiliations**

Stian Soiland-Reyes <u>https://orcid.org/0000-0001-9842-9718</u> Department of Computer Science, The University of Manchester, Manchester, UK; Informatics Institute, University of Amsterdam, NL Genís Bayarri https://orcid.org/0000-0003-0513-0288

Institute for Research in Biomedicine (IRB Barcelona), The Barcelona Institute of Science and Technology (BIST), Barcelona, Spain

Pau Andrio https://orcid.org/0000-0003-2116-3880

Barcelona Supercomputing Center (BSC), Barcelona, Spain Robin Long https://orcid.org/0000-0003-2249-645X

Research IT, The University of Manchester, Manchester, UK Douglas Lowe https://orcid.org/0000-0002-1248-3594

Research IT, The University of Manchester, Manchester, UK Ania Niewielska https://orcid.org/0000-0003-0989-3389

European Bioinformatics Institute (EMBL-EBI), Cambridge, UK Adam Hospital https://orcid.org/0000-0002-8291-8071

Institute for Research in Biomedicine (IRB Barcelona), The Barcelona Institute of Science and Technology (BIST), Barcelona, Spain

Paul Groth https://orcid.org/0000-0003-0183-6910

Informatics Institute, University of Amsterdam, NL

# References

- [Andrio 2019] Pau Andrio, Adam Hospital, Javier Conejero, Luis Jordá, Marc Del Pino, Laia Codo, Stian Soiland-Reyes, Carole Goble, Daniele Lezzi, Rosa M. Badia, Modesto Orozco, Josep Ll. Gelpi (2019): BioExcel Building Blocks, a software library for interoperable biomolecular simulation workflows. Scientific Data 6:169 <u>https://doi.org/10.1038/s41597-019-0177-4</u>
- [Beg 2021] Marijan Beg, Juliette Taka, Thomas Kluyver, Alexander Konovalov, Min Ragan-Kelley, Nicolas M. Thiery, Hans Fangohr (2021): Using Jupyter for Reproducible Scientific Workflows. Computing in Science & Engineering 23(2), pp 36–46. https://doi.org/10.1109/mcse.2021.3052101
- [Biomoby 2008] The BioMoby Consortium (2008): Interoperability with Moby 1.0—It's better than sharing your toothbrush! *Briefings in Bioinformatics* 9(3), pp 220–231. <u>https://doi.org/10.1093/bib/bbn003</u>
- [Blankenberg 2014] Daniel Blankenberg, Gregory Von Kuster, Emil Bouvier, Dannon Baker, Enis Afgan, Nicholas Stoler, James Taylor, Anton Nekrutenko, the Galaxy Team (2014): Dissemination of scientific software with Galaxy ToolShed. Genome Biology 15:403 <u>https://doi.org/10.1186/gb4161</u>
- [Brack 2021] Paul Brack, Peter Crowther, Stian Soiland-Reyes, Stuart Owen, Douglas Lowe, Alan R Williams, Quentin Groom, Mathias Dillen, Frederik Coppens, Björn Grüning, Ignacio Eguinoa, Phil Ewels, Carole Goble (2021): 10 Simple Rules for making a software tool workflow-ready. (*submitted*)
- [Crusoe 2021] Michael R. Crusoe, Sanne Abeln, Alexandru Iosup, Peter Amstutz, John Chilton, Nebojša Tijanić, Hervé Ménager, Stian Soiland-Reyes, Carole Goble, The CWL Community (2021): Methods Included: Standardizing Computational Reuse and Portability with the Common Workflow Language. Communication of the ACM (accepted). arXiv:2105.07028 https://doi.org/10.1145/3486897
- [Cohen-Boulakia 2017] Sarah Cohen-Boulakia, Khalid Belhajjame, Olivier Collin, Jérôme Chopard, Christine Froidevaux, Alban Gaignard, Konrad Hinsen, Pierre Larmande, Yvan Le Bras, Frédéric Lemoine, Fabien Mareuil, Hervé Ménager, Christophe Pradal, Christophe Blanchet (2017): Scientific Workflows for Computational Reproducibility in the Life Sciences: Status, Challenges and Opportunities. Future Generation Computer Systems 75, pp 284–298. https://doi.org/10.1016/j.future.2017.01.012
- [De Smedt 2020] Koenraad De Smedt, Dimitris Koureas, Peter Wittenburg (2020): FAIR Digital Objects for Science: From Data Pieces to Actionable Knowledge Units. Publications 8(2):21 <u>https://doi.org/10.3390/publications8020021</u>
- [Garijo 2011] Daniel Garijo, Yolanda Gil (2011): A New Approach for Publishing Workflows. Proceedings of the 6th Workshop on Workflows in Support of Large-Scale Science - WORKS '11. <u>https://doi.org/10.1145/2110497.2110504</u>
- [Garijo 2014] Daniel Garijo, Pinar Alper, Khalid Belhajjame, Oscar Corcho, Yolanda Gil, Carole Goble (2014): Common Motifs in Scientific Workflows: An Empirical Analysis. Future Generation Computer Systems 36 pp 338–51. https://doi.org/10.1016/j.future.2013.09.018
- [Giovanni 2016] Renato De Giovanni, Alan R. Williams, Vera Hernández Ernst, Robert Kulawik, Francisco Quevedo Fernandez, Alex R. Hardisty (2016): ENM Components: a new set of web service-based workflow components for ecological niche modelling. Ecography 39, pp 376–383. <u>https://doi.org/10.1111/ecog.01552</u>
- [Gray 2017] Alasdair Gray, Carole Goble, Rafael Jimenez, Bioschemas Community (2017): **Bioschemas: From Potato** Salad to Protein Annotation. Poster, *International Semantic Web Conference (ISWC)*, Vienna Austria, 2017-10-23. https://iswc2017.semanticweb.org/paper-579/

- [Goble 2020] Carole Goble, Sarah Cohen-Boulakia, Stian Soiland-Reyes, Daniel Garijo, Yolanda Gil, Michael R. Crusoe, Kristian Peters, Daniel Schober (2020): FAIR Computational Workflows. Data Intelligence 2(1), pp 108–121. https://doi.org/10.1162/dint a 00033
- [Grüning 2018a] Björn Grüning, John Chilton, Johannes Köster, Ryan Dale, Nicola Soranzo, Marius van den Beek, Jeremy Goecks, Rolf Backofen, Anton Nekrutenko, James Taylor (2018): Practical Computational Reproducibility in the Life Sciences. Cell Systems 6(6) pp 631–635. <u>https://doi.org/10.1016/j.cels.2018.03.014</u>
- [Grüning 2018b] Björn Grüning, Ryan Dale, Andreas Sjödin, Brad A. Chapman, Jillian Rowe, Christopher H. Tomkins-Tinch, Renan Valieris, the Bioconda Team, Johannes Köster (2018): **Bioconda: Sustainable and Comprehensive Software Distribution for the Life Sciences**. *Nature Methods* **15**, pp 475–476. https://doi.org/10.1038/s41592-018-0046-7.
- [Hardisty 2021] Alex Hardisty, Paul Brack, Carole Goble, Laurence Livermore, Ben Scott, Quentin Groom, Stian Soiland-Reyes (2021): The Specimen Data Refinery: A canonical workflow framework and FAIR Digital Object approach to speeding up digital mobilisation of natural science collections. Submitted, *Data Intelligence, this issue.*
- [Hospital 2020] Adam Hospital, Genís Bayarri, Stian Soiland-Reyes, Jose Lluis Gelpi, Pau Andrio, Daniele Lezzi, Sarah Butcher, Ania Niewielska, Yvonne Westermaier, Rosa Maria Badia, Rodrigo Vargas, Alexandre Bonvin (2020): BioExcel-2 Deliverable 2.3 – First release of demonstration workflows. Project deliverable, Zenodo. https://doi.org/10.5281/zenodo.4540432
- [Ison 2021] Jon Ison, Hans Ienasescu, Emil Rydza, Piotr Chmura, Kristoffer Rapacki, Alban Gaignard, Veit Schwämmle, Jacques van Helden, Matúš Kalaš, Hervé Ménager (2021): biotoolsSchema: a formalized schema for bioinformatics software description. GigaScience, 10(1):giaa157 <u>https://doi.org/10.1093/gigascience/giaa157</u>
- [Jupyter 2018] Jupyter Project, Matthias Bussonnier, Jessica Forde, Jeremy Freeman, Brian Granger, Tim Head, Chris Holdgraf, et al. (2018): Binder 2.0 - Reproducible, Interactive, Sharable Environments for Science at Scale. Proceedings of the 17th Python in Science Conference. SciPy, 2018. <u>https://doi.org/10.25080/majora-4af1f417-011</u>
- [Katz 2021] Daniel S. Katz, Morane Gruenpeter, Tom Honeyman, Lorraine Hwang, Mark D. Wilkinson, Vanessa Sochat, Hartwig Anzt, Carole Goble, FAIR4RS Subgroup 1 (2021): A Fresh Look at FAIR for Research Software. arXiv:2101.10883
- [Lamprecht 2020] Anna-Lena Lamprecht, Leyla Garcia, Mateusz Kuzak, Carlos Martinez, Ricardo Arcila, Eva Martin Del Pico, Victoria Dominguez Del Angel, et al. Towards FAIR Principles for Research Software. Data Science 3(1), pp 37–59. <u>https://doi.org/10.3233/ds-190026</u>
- [Leipzig 2020] Jeremy Leipzig, Daniel Nüst, Charles Tapley Hoyt, Stian Soiland-Reyes, Karthik Ram, Jane Greenberg (2020): The role of metadata in reproducible computational research. arXiv:2006.08589
- [McMurry 2017] Julie A McMurry, Nick Juty, Niklas Blomberg, Tony Burdett, Tom Conlin, Nathalie Conte, Mélanie Courtot, John Deck, Michel Dumontier, Donal K Fellows, Alejandra Gonzalez-Beltran, Philipp Gormanns, Jeffrey Grethe, Janna Hastings, Jean-Karim Hériché, Henning Hermjakob, Jon C Ison, Rafael C Jimenez, Simon Jupp, John Kunze, Camille Laibe, Nicolas Le Novère, James Malone, Maria Jesus Martin, Johanna R McEntyre, Chris Morris, Juha Muilu, Wolfgang Müller, Philippe Rocca-Serra, Susanna-Assunta Sansone, Murat Sariyar, Jacky L Snoep, Stian Soiland-Reyes, Natalie J Stanford, Neil Swainston, Nicole Washington, Alan R Williams, Sarala M Wimalaratne, Lilly M Winfree, Katherine Wolstencroft, Carole Goble, Cristopher J Mungall, Melissa A Haendel, Helen Parkinson (2017): Identifiers for the 21st century: How to design, provision, and reuse identifiers to maximize utility and impact of life science data. *PLOS Biology* 15(6):e2001414 https://doi.org/10.1371/journal.pbio.2001414
- [Möller 2017] Steffen Möller, Stuart W. Prescott, Lars Wirzenius, Petter Reinholdtsen, Brad Chapman, Pjotr Prins, Stian Soiland-Reyes, Fabian Klötzl, Andrea Bagnacani, Matúš Kalaš, Andreas Tille, Michael R. Crusoe (2017): Robust cross-platform workflows: How technical and scientific communities collaborate to develop, test and share best practices for data analysis. Data Science and Engineering 2, pp 232–244. <u>https://doi.org/10.1007/s41019-017-0050-4</u>
- [Saltz 2006] Joel Saltz, Scott Oster, Shannon Hastings, Stephen Langella, Tahsin Kurc, William Sanchez, Manav Kher, Arumani Manisundaram, Krishnakant Shanbhag, Peter Covitz (2006): caGrid: design and implementation of the core architecture of the cancer biomedical informatics grid. *Bioinformatics* 22(15), pp 1910–1916, https://doi.org/10.1093/bioinformatics/btl272
- [Soiland-Reyes 2021] Stian Soiland-Reyes, Peter Sefton, Mercè Crosas, Leyla Jael Castro, Frederik Coppens, José M. Fernández, Daniel Garijo, Björn Grüning, Marco La Rosa, Simone Leo, Eoghan Ó Carragáin, Marc Portier, Ana Trisovic, RO-Crate Community, Paul Groth, Carole Goble (2021): Packaging research artefacts with RO-Crate. Data Science (accepted). arXiv:2108.06503v1
- [Stodden 2016] Victoria Stodden, Marcia McNutt, David H. Bailey, Ewa Deelman, Yolanda Gil, Brooks Hanson, Michael A. Heroux, John P.A. Ioannidis, Michela Taufer (2016): Enhancing reproducibility for computational methods. *Science* 354(6317), pp 1240–1241 <u>https://doi.org/10.1126/science.aah6168</u>
- [Wittenburg 2021] Peter Wittenburg et al. (2021): CWFR Position Paper. OSF, January 6, 2021. https://osf.io/3rekv/