Phenotype Ontologies Traversing All The Organisms (POTATO) workshop: 2nd edition

Short report

Nicolas Matentzoglu[1,2], James P. Balhoff[2,3], Susan M. Bello[4], Yvonne M. Bradford[5,6], Leigh C. Carmody[2,7], Laurel D Cooper[8], Virginie Courtier-Orgogozo[9], Alayne Cuzick[10], Wasila M. Dahdul[11], Alexander D. Diehl[12], Stacia Engel, Petra Fey, Malcom Fisher[13], Christian A. Grove[14,15], Melissa A. Haendel[2,16,17], Midori A. Harris[18], Nomi L. Harris[2,19], Sebastian Köhler[2,20,21], Marie-Angélique Laporte[22], Eunice Y. McMurray[23], Julie A. McMurry[2,16], Annalisa Milano[1], Chris Mungall[2,19], Monica C. Munoz-Torres[2,16], Clare Pilgrim[23,24], Ajay Pillai, Sofia MC Robb[25], Peter N. Robinson[2,7,26], Paul Schofield, Erik Segerdell[13,27], Achchuthan Shanmugasundram, Sabrina Toro[6], Nicole Vasilevsky[2,17], Shur-Jen Wang, Val Wood[18], David Osumi-Sutherland[1]

1 European Bioinformatics Institute (EMBL-EBI), Wellcome Trust Genome Campus, Cambridge, UK; 2 Monarch Initiative, monarchinitiative.org; 3 Renaissance Computing Institute, University of North Carolina at Chapel Hill; 4 The Jackson Laboratory, Bar Harbor, ME, USA; 5 University of Oregon, Institute of Neuroscience; 6 Zebrafish Information Network, University of Oregon, Eugene, OR 97403; 7 The Jackson Laboratory for Genomic Medicine, Farmington CT 06032, USA; 8 Department of Botany and Plant Pathology, Oregon State University; 9 CNRS, Institut Jacques Monod; 10 Rothamsted Research; 11 University of South Dakota; 12 Department of Biomedical Informatics, Jacobs School of Medicine and Biomedical Sciences, University at Buffalo, Buffalo, NY 14203; 13 Xenbase, Cincinnati Children's Hospital Medical Center; 14 California Institute of Technology; 15 Wormbase, www.wormbase.org; 16 Linus Pauling institute, Oregon State University, Corvallis OR, USA; 17 Oregon Health & Science University, Portland, OR 97217; 18 PomBase, Department of Biochemistry and Cambridge Systems Biology Centre, University of Cambridge, Cambridge, UK; 19 Environmental Genomics and Systems Biology, Lawrence Berkeley National Laboratory, Berkeley, CA, 94720, USA; 20 Charité Centrum für Therapieforschung, Charité - Universitätsmedizin Berlin Corporate Member of Freie Universität Berlin, Humboldt-Universität zu Berlin, and Berlin Institute of Health, Berlin, 10117, Germany; 21 Einstein Center Digital Future, 10117, Berlin, Germany; 22 Bioversity International; 23 University of Cambridge; 24 FlyBase, Department of Genetics, University of Cambridge, Downing Street, Cambridge, UK; 25 Stowers Institute for Medical Research, Kansas City, MO ; 26 Institute for Systems Genomics, University of Connecticut, Farmington, CT, USA; 27 Institute of Ecology and Evolution, University of Oregon, Eugene, OR

Introduction and Motivation for the Workshop

The Phenotypes Traversing All the Organisms Workshop (POTATO) series brings together phenotype ontology curators and developers to reconcile phenotype ontologies across species. In the first workshop (Oregon, 2018), we introduced methods and tools for community-driven development of logical definitions for phenotypes. Together we resolved some of the logically divergent definitions across organisms. As a result, 14 phenotype ontologies and databases covering all major model organisms joined a common Phenotype Ontology Reconciliation Effort[1] with bi-weekly meetings and focus groups. A report of the first installment of the workshop, including a detailed description outlining the motivation of the workshop series, is available on Zenodo (1). The central purpose of this workshop was to bring together users and developers of phenotype ontologies and to provide them with a forum to share, learn and debate. The second edition of the workshop further aimed to address two of the fundamental bottlenecks identified in the course of these community efforts:

[1] https://github.com/obophenotype/upheno/wiki/Phenotype-Ontologies-Reconciliation-Effort

- The Phenotype And Trait Ontology (PATO), an essential driver of inference in phenotype ontologies, has limitations that can make it hard to understand and use in defining phenotypes. For example, textual definitions need to be added and improved, and hierarchies refined to be more intuitive.
- The phenotype ontology community needs efficient ways to directly contribute to anatomical and cell concepts in the species-independent anatomy ontology Uberon and the Cell Ontology (CL).

Workshop summary

The 2nd edition of the POTATO workshop was co-located with Biocuration 2019[2] and brought together 24 curators and ontology developers from a variety of backgrounds. As can be seen below, roughly half of the participants have moderate or more experience with ontology development, and more than two thirds have moderate or more experience with using ontologies for curation. More than half of the participants have used PATO occasionally or regularly before.

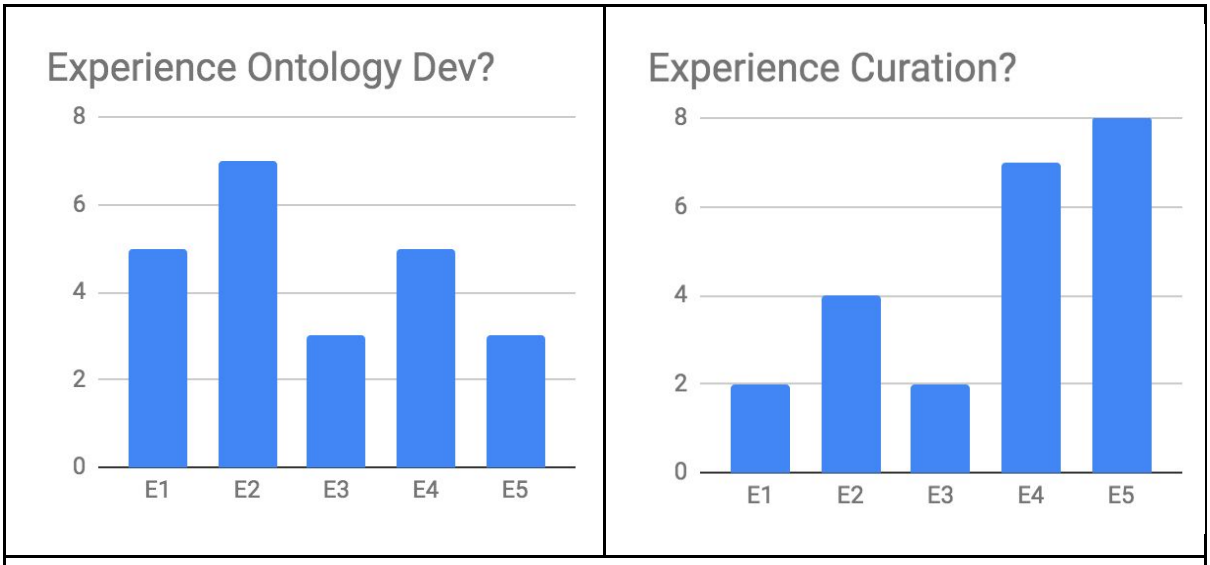

Experience levels of participants on a Likert scale (1-5). E1 indicates very low (novice), and E5 expert-level experience. Y-axis is number of participants.

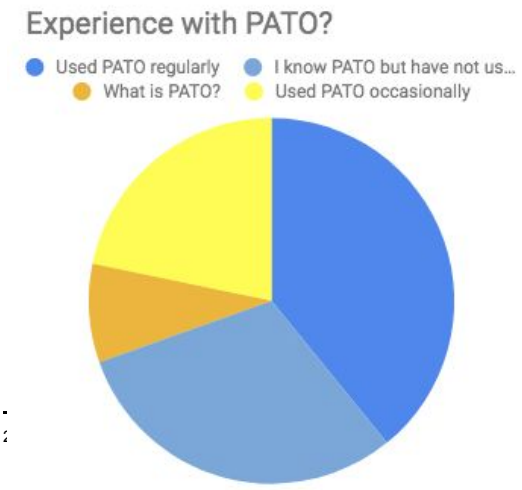

The participants were members of important groups in the phenotype curation space such as the Monarch Initiative, the Alliance of Genome Resources, ZFIN, PomBase, dictyBase, PHIBase, GO, SGD, NIH, HPO, FlyBase, MGI and more.

The main concern for this installment of the workshop was to develop a strategy to deal with shortcomings and current limitations of PATO that were identified by the Phenotype Ontology

Reconciliation Effort during and since the first edition of the workshop. Some of the problem areas were:

- An unclear scope and ontological description of important phenotypic modifiers such as ‘abnormal’, ‘normal’ and ‘pathological’.
- A lack of clarity in many of the textual definitions that were originally geared towards formal and philosophically minded people rather than biological curators, and a lack of textual definitions in general.
- A lack of clarity of when to use key concepts around processual frequency, in particular rate vs occurrence vs frequency.
- The unclear scope of one of PATO’s central branches: morphology.
- The problem of ‘increased amount’ vs ‘has extra parts of type’ when talking about increased counts of anatomical entities.
- The problem of how to faithfully model ‘absence’ when logically defining phenotypes.

The workshop started out with a few general introductions of the reconciliation effort, achievements to date and a general introduction to the rich history of PATO and its primary grouping classes. We then divided the participants into 5 focus groups and introduced them to 7 problems we had selected beforehand (details in Section “Key outcomes”). The focus groups were instructed to discuss the problems and develop solutions. Two 60-90 min discussion sessions were followed by an extensive debate of the observations made with the whole group. During the second part of the workshop, participants were introduced to ontology development workflows with Git. Again within their focus groups, they made changes to PATO with straightforward fixes for issues they had identified during the discussion session (and documented as GitHub issues), created pull requests directly on the public PATO repository and reviewed each other’s pull requests. In the third and last part of the workshop, the participants were introduced to Uberon editing workflows. Uberon is a widely-used, species-independent anatomy ontology that was recently opened to community editing.

Key Outcomes

New editing workflows for the species-independent anatomy ontology Uberon and the Cell Ontology

The species-independent anatomy ontology Uberon was, just prior to the workshop, officially lifted to be a community-editable resource. Uberon is officially hosted at https://github.com/obophenotype/uberon. All editing of Uberon happens through pull requests to ensure that any edits made are checked by quality control prior to inclusion in the ontology, and to give other members of the community a chance to comment on the change. Together with PATO and GO, Uberon is the most important reference ontology driving phenotypic integration. The underlying hypothesis is that phenotypes affecting the same or homologous anatomical locations often share some underlying genetic causes. For example, increased paw size in mice and increased hand size in humans are likely to at least share some underlying genetic mechanisms.

The challenge for the community over the next years will be to evolve Uberon to a truly species-independent ontology covering all Metazoa, and to ensure that it integrates seamlessly with ontologies covering other relevant taxa such as plants, yeasts and procaryotes. While Uberon and CL conceptually cover all Metazoa, their current axiomatization is centered around vertebrates. Moreover, it is likely that many organism communities such as Zebrafish and C. elegans will continue to use their own anatomy ontologies. In such cases, phenotypic integration will rely on aligning species-specific phenotype ontologies with Uberon or CL as an upper, species-independent layer. Alignment, in this case, means to link a species-specific concept, such as Zebrafish Anatomy Ontology's ZFA:Eye, with the corresponding species-independent concept, such as UBERON:Eye, typically using a subclass-of axiom or (as is currently the case) by logically defining ZFA:Eye in Uberon as 'Uberon:eye and part-of some Danio rerio'. This will not only involve the inclusion of new classes of anatomical entities, but also the careful revision of existing general assumptions. For that reason, the Uberon development team decided to open Uberon up to the community. To accommodate diverging interests across a large number of current and prospective users, an editing workflow is proposed that allows members from across the community to comment on, accept and reject proposed changes:

1. Create an issue with the proposed change (clear title, and detailed description) and give reasons why. This step can be omitted in cases of extreme urgency, but helps with general transparency. Clearly indicate whether you plan to make the change yourself.
2. Ask for edit rights on the Uberon issue tracker in a clearly labelled, separate issue
3. Edit https://github.com/obophenotype/uberon/blob/master/uberon_edit.obo[3] and/or https://github.com/obophenotype/cell-ontology/blob/master/src/ontology/cl-edit.owl using Protege, save and commit the change **to a branch.**
4. Create a pull request. Link any existing issues, and re-iterate the rationale and nature for the change. Identify suitable reviewers if possible.
5. Uberon's core editor team will give the community 7 days to review your change. If no one vetoes your change, they will accept your pull request and merge it into master. If 7 days have gone by with no response, you can chase the core editing team with a comment on your pull request.

For the foreseeable future, it is likely that organism-specific communities will continue to maintain their own anatomy ontologies, which will be linked up to Uberon to facilitate phenotypic integration. An exception should be anything that relates to the taxon of Mammalia: both the Mammalian and Human Phenotype Ontologies directly reference Uberon terms in their logical definitions. To maintain conceptual consistency between species-specific anatomy ontologies (SSPO) and Uberon, we will likely have to develop a way to 'push down' modelling suggestions from Uberon to the SSPOs, whenever fundamental modelling inconsistencies are uncovered (for example, is an anatomical line a material entity?; which relations are used to denote developmental predecessors?).

[3] Xrefs to other anatomy ontologies are also maintained in the uberon_edit.obo file

The history of PATO

PATO was started in 2002 by Michael Ashburner and George Gkoutos, with the goal of semantically representing phenotypic qualities, including properties, attributes and characteristics (2). There are approximately 1700 classes in PATO that are used across 20 phenotype ontologies in about 120,000 axioms. In total, PATO is used in 63 ontologies, in about 185,000 axioms. It is widely used for direct annotation of phenotypes, in post-compositional entity-quality representations with anatomical or other ontology terms (such as quality: PATO_0002359 broad, entity: UBERON_0000004 nose). PATO is a central component of the phenotype ontology alignment and uPheno efforts.

At the workshop, we aimed to improve the usability and utility of PATO by obtaining feedback from the workshop participants, enabling community contributions by training participants on editing the ontology and performing GitHub pull requests, and thereby improving inference of classification in phenotype ontologies.

Strategy for addressing fundamental PATO issues

At the 2nd POTATO workshop, participants broke out into small groups and worked on a set of exercises, outlined below and available here.

Normal vs abnormal

PATO currently has this class:

deviation (from normal)
← **abnormal**: "... deviation from normal *or* average."
← **normal**: "... no deviation from normal *or* average."

All workshop participants agreed that the location of 'normal' under deviation from normal made no sense. We therefore edited PATO to move 'normal' to be a sibling of 'deviation (from normal)'. There was also general consensus that including the clause 'or average' in the definitions does not make sense as this would make almost everything abnormal; we will therefore remove this.

This still begs the question of how and when we should distinguish normal from abnormal. Discussion of this topic resulted in some lively debate.

Evolutionary biologists (represented by Phenoscape (3) and Gephebase curators) made it clear that this distinction is not relevant for describing naturally occurring phenotypes from the comparative, evolutionary literature. In this context, precisely described variation between taxa or between natural populations is important for curating phenotypic information (e.g. in a character matrix). Model organism biologists and clinicians can agree that many phenotypes are abnormal and find it useful to distinguish these from normal phenotypes (corresponding to wild-type, canonical or non-pathological phenotypes). For example, in humans, blue eyes is a normal phenotype, but fused eyes is an abnormal phenotype. Model organism biologists and clinicians mostly care about abnormal phenotypes. This is reflected

in the structure of phenotype ontologies, which frequently assert 'abnormal phenotype' to be the root of their ontologies (HP) or of all branches (MP).

There was a lively discussion of the relationship of the terms normal and abnormal to the concept of a wild-type. Attendees also made the point that sometimes it is hard to tell when curating whether a phenotype is abnormal compared to wild-type or just differs from the control presented in the paper being curated. One possible solution to this is to encourage the use of phenotype terms that are agnostic with respect to normality/abnormality. This can easily be supported by PATO and logical definition design patterns. We agreed to provide these patterns, leaving the choice of whether to create terms using them to the individual phenotype ontologies.

Outcomes:

1. Edits to the PATO hierarchy and definitions for normal and abnormal have removed obvious errors.
2. We have not agreed on a complete definition of abnormal, will consider adding a usage comment that enumerates at least some usage.
3. We will provide design patterns to give phenotype ontology developers the option of including terms that are agnostic regarding normality.

Abnormal vs pathological

In PATO, pathological is currently a subclass of abnormal (meaning that all pathological phenotypes are inferred to be abnormal) and this is reflected in the definition: "... abnormal and having a destructive effect on living tissue."

The focus groups were asked to answer the question: Are pathological phenotypes always abnormal? If the answer is no, do we have use cases for pathological phenotypes that are not abnormal?

The general consensus was that pathological phenotypes can sometimes be normal. For example, some phenotypes are considered pathological in younger people but may not be pathological in older adults (like vision changes, memory loss, etc.). Inflammatory responses such as fever are normal processes; however, they turn pathological if they are dysregulated. This suggests that pathological should not remain a subclass of abnormal.

For a classic discussion of the problem of setting the boundary between normal and pathological see (4).

Revising PATO textual definitions

Question: PATO textual definitions typically have the structure: An X quality inhering in a bearer by virtue of the bearer's (disposition to) Y. How can we simplify these to make them more accessible?

General consensus: We should delete some of the preamble words to make definitions more succinct. Attendees gave some examples of how this could work:

- Amorphous: “A morphology quality inhering in a bearer by virtue of the bearer's lack of distinct morphology.” -> "Lack of distinct morphology"
- Contractility: "A physical quality in which the entity has the ability to shrink or contract." -> "The ability to shrink or contract."

However, Chris Mungall suggested that including a genus (superclass) in definitions is helpful for ontology editors. For example:

> A physical quality that is a lack of distinct morphology
>
> A physical quality that is the ability to shrink or contract

How should we improve definitions and guidance about when to use *rate* vs *occurrence* vs *frequency*?

PATO has terms for rate, frequency and occurence. Each of these has subclasses that can be used to record when rate frequency or occurrence of a process are increased or decreased. Curators and ontologists using these terms have reported that it can be hard to know which of these terms to choose, given their current definitions:

physical quality of a process (18341)

> **rate (4191):** "A quality of a single process inhering in a bearer by virtue of the bearer's occurrence per unit time."
>
> > **frequency (356):** "A physical quality which inheres in a bearer by virtue of the number of the bearer's repetitive actions in a particular time."

temporal distribution quality (11390): "A temporal distribution pattern of process occurrences within a regulation/reference process."

> **occurrence (9589):** "A quality of a single process inhering in a bearer by virtue of the bearer's occurrence."

The following recommendations were made based on discussion:

We should distinguish two types of rate: The rate of regular occurrence of discrete events (e.g. heart rate is the rate of occurrence of heart beats), vs the rate of occurrence of some continuous process such as growth. Frequency is a subclass of the former, not the latter. Frequency should also be a subclass of occurrence - covering regular, repetitive occurrence. Occurrence itself covers both regular and irregularly occurring events. For example, fly grooming does not occur at regular, predictable intervals but behavioral phenotypes can affect how often it takes place. We have begun to implement these recommendations, adding new terms to PATO, modifying existing definitions and adding comments on usage.

"Increased amount" vs "has extra parts of type"

PATO has both 'increased amount' and 'has extra parts of type'. Do we need both? If so, how can we improve definitions or add usage statements to ensure consistent usage?

'has extra parts of type' can be used to record increased counts in a specific location, while 'increased amount' cannot, so ‘has extra parts of type’ might be preferable in such cases. However, it also requires logical definition structures that make it hard to automate classification of increased counts of X under other X phenotypes.

Increased amount can be usefully applied to processes with measurable quantitative outcomes (continuous variables), e.g. growth, whereas extra parts applies to counts (discrete variables). One possibility would be to use the amount branch for processes only.

How do we represent *absence* and *reduced counts*?

Modelling of absence is a tricky ontological problem (5). While acknowledging that the problem was likely too complex to solve in a short workshop session, we asked attendees to address a small, focussed set of questions in order to guide further discussion. We have since followed this up with further discussion in our biweekly meetings. Some content here reflects those discussions.

General discussion points: Biologists record absence phenotypes in relation to what is expected/normal for a species (model organism biologists and physicians) or in the context of a comparison between related species (evolutionary biologists). A physician wouldn't record that their patient has an 'absent tail' phenotype, even though this is correct from a strict logical perspective. In contrast, absence of a tail in humans is interesting from an evolutionary perspective.

Focussed questions: **(1) Should absent be a subclass of morphology?** Having this structure in PATO will lead for example to the inference that (abnormally) absent teeth is a subclass of abnormal tooth morphology. While some phenotype ontologies (MP and HP) already assert this manually, multiple focus groups objected that this asserts morphological phenotypes for the wrong structures. Absent teeth is not a tooth morphology phenotype but it is a phenotype of the mouth and of dentition. As a result, we have not changed the PATO hierarchy, but will investigate whether we might be able to infer morphology phenotypes from phenotypes that result in the absence of parts. **(2) Should absent be a subclass of decreased amount?** Attendees were happier with this assertion. It is reflected in manually asserted classifications in MP and HP: e.g. MP has 'absent teeth' under 'reduced teeth numbers'. For discrete entities that are present normally with (normally) a count of one (head, tail), some attendees were a little more uncomfortable with this: Calling no-tail a subclass of decreased amount of tails would be odd. On the other hand, this may not be an issue as long as the 'decreased number' classes are not present, which seems unlikely for tail or head. **(3) Should we use PATO:'absent' to record absence phenotypes or PATO:'lacks all parts of type'?** The latter sits under PATO classes used to report reduced number (and so can be used to drive the inference we agreed was useful in discussion of question 2). It is also used in logical definition patterns that allow local absence to be recorded (e.g. absence of hair on the head, rather than an absence of all hair). These considerations suggest we should favour 'lacks all parts of type'. However, further discussion of logical consequences is needed before a final decision is made.

What is the scope of PATO's *morphology*?

Morphology and its subterms are the most widely used branch of PATO outside the very abstract upper levels.

A classical view of biology divides it into morphology (structure) and physiology (function, process). Mouse and Human phenotype ontologies take this broad view - attempting to divide phenotypes cleanly between morphological and physiological phenotypes. In doing so, they include 'color' and 'composition' down to the molecular level under morphology. A more restricted use (followed by FYPO and GO) limits morphology to refer to just 'size' and 'shape' (consistent with the latin roots of the term (study of shape)).

PATO has a relatively broad view of what counts as morphology: "A quality of a single physical entity inhering in the bearer by virtue of the bearer's size or shape or structure." This definition does not cover color. The major (most used) classes under morphology reflect this definition well:

physical entity quality (119881)
- ← **morphology** (54610)
  - ← **size** (16395)
  - ← **shape** (7969)
  - ← **structure** (8967)
    - ← **composition** (1840)

Discussion revolved around whether color and molecular composition should be included under morphology. The color of a living entity is the result of its surface properties, its transmission and emission properties. Coloration can thus be divided into chemical color (pigmentation) and structural color (such as iridescence, which is created by microscopic structures of surfaces). The group reached a consensus that color should be included under morphology (although some had reservations). As a result, we have since edited PATO to reflect this. The subject of whether to include molecular level composition proved more controversial.

This is reflected in PATO:

composition
- ← **amylose composition**
- ← **biomaterial purity**
- ← **calcification**

And in HPO

Abnormal liver morphology
- ← **Abnormal hepatic iron concentration**
- ← **Depletion of mitochondrial DNA in liver**

The main concern expressed was that including molecular-level composition makes 'morphology' so broad that it is likely to encompass all phenotypes. Any physiological (process/function) phenotype will almost certainly have accompanying molecular level effects, such as changes in gene expression or metabolites. We reached no firm conclusion on how to resolve this, but plan to discuss it further with the aim of either tightening the definition or coming up with guidance for what types of compositional change should count as morphological.

Outcome summary: color is now a subclass of morphology in PATO. A decision on the placement of molecular composition phenotypes requires further discussion.

Conclusions and outlook

The central objective of the Phenotype Ontology Reconciliation Effort is to enable the scalable implementation of logical descriptions for phenotype terms that are (1) internally consistent and fully amenable to automated classification methods based on OWL reasoning and (2) externally consistent and amenable to cross-species reasoning based on semantic similarity approaches. This objective will be reached through three central (partially overlapping) sub-goals:

1) Developing formal templates for the representation of phenotypes and implementing logical descriptions as instances of these patterns (patternization)
2) Reconciling branches in species-specific phenotype ontologies by ensuring that they refer to the same phenotype patterns whenever appropriate.
3) Aligning reference ontologies (such as anatomy ontologies) and ensuring that the *ontological commitment* (i.e. the assumptions about the world expressed by the axioms used) is shared across all members of the reconciliation effort as much as possible.

This workshop made progress mostly on sub-goal 3 - PATO, GO and Uberon are the most important reference ontologies used to define phenotypes of Metazoa. This goal is also likely to be the hardest and most long-term of the three sub-goals, due to historical, terminological and conceptual differences of different research groups.

Since the workshop, we have started using tools that facilitate the formal review of proposed phenotype design patterns[4]. Our goal is to achieve near complete patternization (goal 1) of existing logical descriptions by the end of 2019 (90%), and at least 30% coverage of logical definitions across all phenotype terms in any given phenotype ontology. With regards to goal 3, we aim to finalise our decisions on the PATO issues debated as part of the workshop around autumn 2019. The next installment of the POTATO workshop series is likely to take place in early 2020.

To join the Phenotype Ontology Reconciliation Effort or for any other questions, please contact Nico Matentzoglu: nicolas.matentzoglu@ebi.ac.uk or Nicole Vasilevsky:

[4] https://github.com/obophenotype/upheno/tree/master/src/patterns/

vasilevs@ohsu.edu

Acknowledgements: The workshop was funded by NIH-UDP: HHSN268201300036C. It contributed to the ongoing work of the Phenotype Ontology Reconciliation Effort as part of the Monarch Initiative, NIH Office of the Director Grant #5R24OD011883, as well as by, HHSN268201400093P, NCI/Leidos #15X143. Coffee and other sustenance was kindly provided by the International Society for Biocuration. A big thank you goes to Midori Harris who oversaw the workshop logistics and ensured a surprise-free workshop experience!

References

1. Matentzoglu N, Balhoff JP, Bello SM, Boerkoel CF, Bradford YM, Carmody LC, et al. Phenotype Ontologies Traversing All The Organisms (POTATO) workshop aims to reconcile logical definitions across species [Internet]. Zenodo; 2018 Dec [cited 2019 Jun 4]. Available from: https://zenodo.org/record/2382757#.XPY9XdNKhTY
2. Gkoutos GV, Green EC, Mallon A-M, Hancock JM, Davidson D. Using ontologies to describe mouse phenotypes. Genome Biol. 2005;6(1):R8.
3. Edmunds RC, Su B, Balhoff JP, Eames BF, Dahdul WM, Lapp H, et al. Phenoscape: Identifying Candidate Genes for Evolutionary Phenotypes. Mol Biol Evol. 2016 Jan;33(1):13–24.
4. Joubert J. « Le normal et le pathologique ». Relire Canguilhem. Revue des Sciences Religieuses. 1999;73(4):497–518.
5. Balhoff JP, Dececchi TA, Mabee PM, Lapp H. Presence-absence reasoning for evolutionary phenotypes. arXiv:14103862 [cs, q-bio] [Internet]. 2014 Oct 14 [cited 2019 Jul 14]; Available from: http://arxiv.org/abs/1410.3862