

Are Bio-ontologies Metaphysical Theories?

Oliver M. Lean

Institut supérieur de philosophie, Université catholique de Louvain, 1348, Louvain-la-Neuve, Belgium
omlean@gmail.com

Forthcoming in Synthese

Abstract

Bio-ontologies are digital frameworks for handling biological and biomedical data. They consist of theoretical entities and relations with explicitly defined logical structures and precise definitions, whose purpose is to provide a shared language for representing information to be distributed and integrated across diverse scientific contexts. It is tempting to view bio-ontologies as clear and formal expressions of a scientific community's ontological commitments about their domain of inquiry, and to view their integration as tantamount to the metaphysical unification of science that some philosophers have envisaged. However, I argue that the local, practical, social and technological factors that influence their design prevent us from straightforwardly reading metaphysical conclusions from them. I discuss these complications and suggest how they can be overcome, revealing more general lessons for the development of a well-founded scientific metaphysics.

Acknowledgements

I'm grateful to Amanda Bryant, Steve Elliott, Charles Pence, Luca Rivelli, and anonymous reviewers for invaluable comments on this paper. This work was supported by the Fonds de la Recherche Scientifique - FNRS under grant no. F.4526.19.

1. Introduction

Scientific metaphysics denotes a loose collection of projects which, in one way or another, aim to bring metaphysics in contact with our best or most successful science (Kincaid, 2013)¹. It arose as a reaction to more traditional a priori approaches to metaphysics which, according to their critics, rely too much on armchair methods like appeals to intuition that do not reliably produce true conclusions about reality (Ladyman & Ross, 2007; Bryant, 2020; though see Bennett, 2016). Scientific metaphysics (sometimes called *naturalistic* metaphysics) holds that metaphysics can gain wider intellectual credibility by drawing in principled ways on the sciences themselves—whether on the content of science, its methodology, or both. A basic motivating principle behind scientific metaphysics is that since science investigates reality, and does so with considerable success, studying its success-making features should have implications about what that reality is like in metaphysical terms.

However, there are different views about exactly what the relationship between science and scientific metaphysics is or should be (Chakravartty, 2013). Traditionally, metaphysics has often been taken to be a subject wholly separate from or prior to empirical inquiry. For example, some take metaphysics to be about the basic categories of reality, whereas science only deals with instantiations of those categories (Paul, 2012). Others have argued that metaphysics is prior to science in the sense that it establishes the a priori conceptual framework that is necessary for science to proceed in the first place (Lowe, 2002). Yet the existence of a priori elements in science doesn't imply that those elements cannot themselves change in response to science as it evolves (Friedman, 2001). In any case, the very possibility of a metaphysics informed by the sciences requires that our reasoning about reality can flow from science to metaphysics, not (or not only) in the other direction. Yet without an explicit and principled account of what sort of facts from science might be brought to bear on metaphysical questions, and exactly how and why they are relevant, the very coherence of a scientific metaphysics is on uncertain ground.

To gain traction on these questions, I will consider the implications for scientific metaphysics of a particular and important aspect of contemporary biological science; namely, its increasing reliance on the production and reuse of large quantities of data. These major changes in how science is done have prompted reflection among philosophers and scientists about the promises and challenges they reveal (Searls, 2005; Brigandt, 2013; Leonelli, 2016; Laubichler et al., 2018). The majority of philosophical work has approached this largely from a combination of historical, epistemological and sociological perspectives: Researchers have asked how this new state of affairs has arisen from earlier stages, whether and how data-driven research changes the very nature of scientific knowledge, how social and institutional structures make data-driven knowledge production possible, and so on. In contrast, my focus here is on what the methods and successes of data-driven biology say about scientific metaphysics.

The aspect of data-centric biology I will consider is *bio-ontologies*, which are formal classification systems used by the biological and biomedical sciences to encode, interpret, store, share, discover and integrate the growing volumes of data they produce. (They are a special case of the more general category of ontologies in information science, which covers any computer-based classification system for some domain, scientific or otherwise.) Bio-ontologies occupy a background role in modern science, and rarely gain wider attention in the way that discoveries of new explanations and phenomena do. They might therefore be easily overlooked by philosophers looking to draw on cutting-edge science for metaphysical discussions, yet they are key to contemporary data-centric science and its success.

The goal of the present work is to argue that, for scientific metaphysicians interested in mining contemporary science for insights about metaphysics, bio-ontologies constitute a rich and untapped well of resources. Of course, philosophers engaging in scientific metaphysics (by that name or another) have long been interested in analyzing scientific theories in order to identify their metaphysical aspects or commitments. However, bio-ontologies possess features that make them particularly amenable to this sort of analysis, yet this promise has until now gone largely unrealized. My aim here is therefore metaphilosophical: I will not be appealing to bio-ontologies to argue for a particular metaphysical claim or theory. Rather, I argue that, from the perspective of scientific metaphysics, at least some metaphysical disputes can be usefully addressed by looking at bio-ontologies and the wider scientific context surrounding them, and discuss important factors to consider when doing so.

The reasons for the particular value of bio-ontologies to scientific metaphysics, which I will elaborate throughout, are as follows. First, bio-ontologies offer a particularly explicit and formal expression of a scientific community's current state of biological knowledge. They aim to capture the way scientists currently conceptualize their domain of inquiry in the form of a set of terms and relations with clear logical relationships and careful and precise definitions. This is a result of the demand for data to be handled in computational form in ways that support automation of various uses of that data, as well as the need to incorporate the needs and insights of diverse communities of researchers. Second, a primary function of bio-ontologies is to integrate and synthesise the knowledge of these diverse communities of scientists, with differing methods, domains of inquiry, terms, definitions, research aims, and so on. They are therefore the result of careful and ongoing negotiation between parties about the best way to represent biological knowledge across these differences. Because of this integrative role, bio-ontologies may be particularly relevant to questions about the metaphysical unity of science, since their aim is to unify different scientific disciplines and fields under an agreed conceptualization of the biological world. Third, and finally, these discussions explicitly incorporate conceptual tools developed in analytic metaphysics. That they bear the name "ontology" is no accident: While this derives more directly from computer science, it ultimately has origins in the philosophical sense of the term (Leonelli, 2010)—that is, in the branch of metaphysics concerned with the general categories that constitute reality and how those categories are related. (To distinguish the

two for the purpose of this paper, I will use the capitalized *Ontology* to refer to the branch of philosophy, and lowercase *ontology* or *bio-ontology* to refer to the artifacts used in biological science.) For reasons that will become clear, it is no coincidence that the ongoing design and administration of bio-ontologies has been strongly influenced by the input of metaphysicians and the careful clarifications developed in the discipline.

Taken together, these three reasons offer a strong *prima facie* case for considering bio-ontologies to be of interest to scientific metaphysics. They suggest that bio-ontologies are, in effect, products of the fruitful interaction between the careful conceptualisations of metaphysics on one hand, and empirical science and its concrete practical aims on the other. Le Bihan and Barton (2018) go as far as suggesting that to develop a bio-ontology *is* to build a philosophical Ontology of that domain:

“As a matter of fact, developing rigorous and exhaustive applied ontologies raises issues very similar to those that are classically raised by philosophical ontology, and it has become clear that the success of the former will depend on building upon the theories and methods of the latter. In many respects, building a coherent and solid *applied ontology of biomedicine* amounts to building a formalized, *philosophical ontology of biomedicine*.” (Le Bihan & Barton, 2018, p4, emphasis in the original)

Critically, the fact that bio-ontologies are explicitly designed for practicing scientists with concrete empirical goals gives them the all-important connection to empirical success that so interests scientific metaphysicians. If this is true, it is hard to imagine a more ideal scientific source than bio-ontologies to inform questions about Ontology in philosophy, at least on the face of it: If bio-ontologies are literally the Ontological theories chosen by science, one might think that philosophers aiming to draw metaphysical insights from science can do no better than to look at the bio-ontologies scientists use. For metaphysicians interested in the Ontological structure of the biological world—what sorts of things there are and how they relate to each other—there may be little more to do than to look to the entities and relationships in terms of which biological scientists see fit to encode their empirical data.

The present work critically analyzes this proposal. That is, it considers whether, and in what respects, biological ontologies might be used to inform Ontological questions in scientific metaphysics. As I will elaborate, there are reasons to think that the relationship between bio-ontologies and metaphysics is neither simple nor clear. These complications relate in various ways to the fact that, rather than being attempts to describe the ultimate or universal structure of reality—the purview of Ontology in the classical philosophical sense—bio-ontologies are in the first instance artifacts designed to achieve concrete epistemic aims in relatively local contexts. So if bio-ontologies are to be considered relevant to metaphysics at all, there is work to be done to clarify exactly how their metaphysical significance fits with the epistemic, social, practical, and technological factors that drive their design.

I will argue that, rather than being fatal to the prospect of analyzing bio-ontologies for metaphysical purposes, these complicating factors are simply things to take into account when we analyze bio-ontologies from a metaphysical perspective. This discussion yields lessons, I argue, for the broader question of what scientific metaphysics is or might be. Since the *raison d'être* of scientific metaphysics is the principle that science's empirical success underwrites its relevance to metaphysical questions, this implies that this metaphysical relevance is in some way intricately woven with the practical and epistemic aspects that make science successful. An essential task of scientific metaphysics, then, is to understand this complex relationship.

The structure of the paper is as follows. Section 2 outlines the form and function of bio-ontologies and how they are designed and administered. Section 3 discusses the *prima facie* reasons why bio-ontologies might appear extremely valuable for informing discussions in scientific metaphysics—perhaps even that they simply *are* metaphysical theories—and highlights complicating factors that appear to threaten this promise. Section 4 then proposes ways to address these complications that point to a particular understanding of scientific metaphysics in general: Rather than ignoring or bracketing those practical and epistemic motivations, metaphysicians should instead carefully consider the relationship between those factors, the design features that result from them, and, crucially, the domain of inquiry the bio-ontologies are designed to represent. Section 5 then connects these conclusions to broader questions about the nature and purpose of scientific metaphysics in general. Section 6 concludes.

2. The Form and Function of Bio-ontologies

In the biological and biomedical sciences, high-throughput experimental technologies can produce massive quantities of data about biological phenomena. As a result, data are increasingly being reused for research purposes other than those for which they were produced, and even produced with no particular research question in mind. Data in modern science are therefore a tangible and reusable scientific commodity (Leonelli, 2015). Multiple datasets from different sources and in different forms can be integrated to produce new insights about complex systems; for example, to identify candidate targets for drug interventions (Searls, 2005, Berg, 2014) or to generate new insights about network structures (Bechtel 2020). To make all of this possible, data must take a form that can be encoded and labelled digitally, making it possible to store, distribute, process and integrate them quickly, efficiently, and in high volume. Importantly, beyond mere storage and distribution, the sheer volume of data has led much of its *use*—to draw inferences and make decisions—to be outsourced from humans to computers. To reduce (but not eliminate) the risk of cascading errors in these automated operations, the meaning of the terms and relations being used must be made explicit and precise, so that they can be faithfully adhered to in the computational rules.

Bio-ontologies are a key factor in achieving these aims. Intrinsically, they consist of a set of terms and relations corresponding to the entities and empirical relationships that biological

data are taken to tell us about. They are, in short, formal classification systems (Leonelli, 2016) —they aim to represent the theoretical categories used by a community of scientists in a way that allows data to be labelled, stored, shared, processed, and integrated by computer. Scientific investigation has always involved interpreting empirical observations in terms of a set of categories, sorting those observations into boxes with other instances that share certain things in common. Lorraine Daston (2004) calls the classification systems used in science “applied metaphysics”. Modern bioinformatics has simply formalized this perennial feature of empirical inquiry to implement it in modern digital infrastructures.

According to Leonelli, the classification systems that a research community uses are designed in accordance with, and are thus evidence of, the theoretical picture under which that community is operating:

“Researchers who use bio-ontologies for data retrieval implicitly accept, even if they might not be aware of it, the definition of biological entities and processes contained within bio-ontologies at the moment at which they are consulted, which in turn affects how data are used in subsequent research. It is because of this crucial role in expressing available knowledge of biological phenomena, and thus guiding and structuring subsequent research, that bio-ontologies are best regarded as a form of theory.” (Leonelli, 2016, p. 125)

According to Leonelli (2016), bio-ontologies as theories have a particular scientific purpose: namely, they are attempts to coordinate or synthesise the theoretical backgrounds of communities of researchers. A major source of the value of bio-ontologies lies in their ability to connect the terms in one domain of investigation to terms in another, such as those of different scientific disciplines or of communities focused on different model organisms, thus providing a common language to facilitate interdisciplinary collaboration. The theoretical framework of an individual research community is adapted to its own particular experimental methods, objects of study, research aims, and so on. Bio-ontologies often aim to coordinate *between* individual communities, and so they aim for terms and relations that are neutral between the specific contexts of the communities involved.

A widely-used system of bio-ontologies are those of the Open Biological and Biomedical Ontology (OBO), governed by the OBO Foundry (Smith et al., 2007). The stated aim of the OBO Foundry is “to develop a family of interoperable ontologies that are both logically well-formed and scientifically accurate” (<http://obofoundry.org>) — in other words, it aims to develop a growing and evolving family of ontologies that are both comprehensive of their individual domains and well-connected to others in the family. Coordination and integration efforts are thus applied both *within* bio-ontologies—between members of the community that uses a given ontology—and *between* ontologies used by different communities.

One OBO ontology, from which I will draw my examples, is the Gene Ontology (GO). The GO aims to be an authoritative and general-purpose repository for information about the functions of genes (and gene products, also called “genes” as a shorthand) across different biological taxa. Like all OBO ontologies, the GO is essentially a graph-theoretic object consisting of *terms* (nodes) and *relations* between those terms. As we’ll see, these terms and relations are governed by strict rules that determine the logical and inferential implications of the information being represented.

The GO consists of three *sub-ontologies*—molecular functions, cellular components, and biological processes. For example, *mitochondrion* is a *cellular component* (Fig. 1). Together, these three sub-ontologies aim to provide the basis for representing scientific knowledge of the functions of genes across different biological taxa: what the gene products interact with, what processes they are involved in, which part of which organisms those processes occur in, and so on. As an OBO ontology, the GO is subject to the OBO’s strict rules for inclusion. For example, all OBO ontologies must be connected (via relations between terms) to the basic formal ontology (BFO) that provides the foundations on which the terms of the other ontologies are built.

This formalized model of biological processes, their relations and components is structured by the relations between the terms. Some of these are generic relations applicable outside biology, such as *is_a* and *part_of*, while others are more domain-specific, such as *negatively_regulates*. Each of these has its own set of properties, such as transitivity, and rules for their application. For example, the relation *is_a* (“is a subclass of”) must only occur within the sub-ontologies; a biological process cannot be a subclass of a cellular component. The sub-ontologies are connected to each other by other relations; a biological process can bear the *occurs_in* relation to a cellular component, for instance. GO terms are also connected to terms in other BFO ontologies, such as ChEBI (Chemical Entities of Biological Interest) which deals with small chemical compounds.

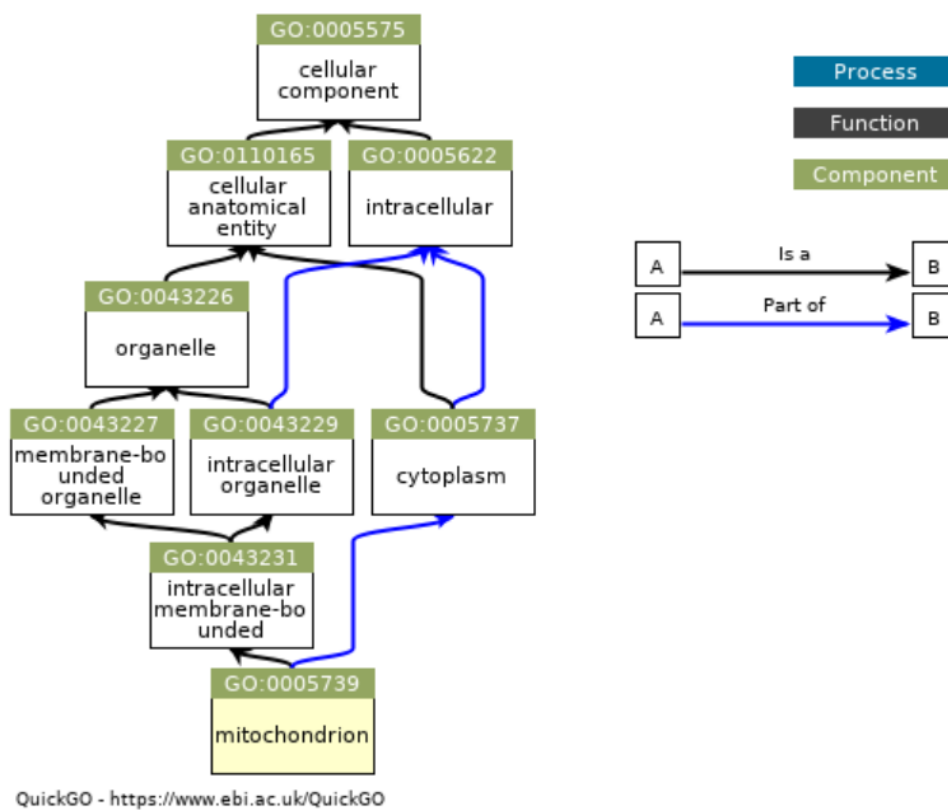


Fig. 1: Ancestor chart for *mitochondrion* in the Gene Ontology. That the mitochondrion is a cellular component is shown by its connection to the root term *cellular component* by a chain of *is_a* relations. It also bears the relation *part_of* to other cellular components, as well as other types of relation to other terms (not shown).

This logical structure allows the GO to serve its purpose of representing empirical information about gene function. This information takes the form of *annotations*, which express claims about the functions of genes by connecting gene names to the terms of the GO via its relations. Each annotation is accompanied by at least one piece of evidence (classified according to the Evidence and Conclusion Ontology, another member of the OBO ontology family), whether direct experimental evidence, or indirect kinds such as inference from similarity with another gene. Many annotations are verified manually by researchers with expertise in both the relevant biology and in formal knowledge representation, often in response to suggestions generated by text mining software that parse scientific literature to extract its empirical claims. Natural language processing tools can more generally be used to interpret scientific texts by connecting its contents to the terms of an ontology (Hirschman et al., 2012; Hoehndorf et al., 2015).

Importantly, many GO annotations are inferred automatically based on the ontology's logical structure. For example, if gene G is annotated as involved in process P, and process P is a

part of process Q, then gene G is inferred to be involved in process Q. The logical axioms that constitute the ontology can therefore be put to service in automated reasoning—to detect contradictions or answer complex queries (Hoehndorf et al., 2015). Ultimately, however, all evidence derives, more or less directly, from explicitly cited sources in the scientific literature. Annotations can be added, removed, and changed in light of new empirical evidence and requests from researchers. The GO thus aims to constitute a “snapshot” of current biological knowledge by being responsive in this way to scientific developments.

Aside from updating annotations that connect genes to the ontology, another way in which the GO responds to scientific feedback is by updating the terms and relations that constitute the ontology—the molecular functions, biological processes and cellular components it includes, how they are defined, and the relations between them. Updates to the ontology are performed in response to user requests and through collaboration with specialist curators. For example, as Leonelli et al. (2011) describe, the term “serotonin secretion” was revised when it was noticed that it erroneously implied (via *is_a* relations) that “serotonin secretion during acute inflammatory response” was an instance of “neurotransmitter secretion”. In acute inflammatory response serotonin’s targets are not neurons, so a new tree structure was generated to reflect this. In general, the evolution of the GO and other bio-ontologies involves careful discussion of the definitions and logical relationships between its terms (Leonelli et al. 2011; Leonelli 2010). As well as cultivating internally coherent and scientifically sound logical structure, good design and curation of bio-ontologies also requires understanding the relationship between its terms and the different theories, purposes, and experimental methods of the scientific disciplines involved (Leonelli, 2010; Smith & Ceusters 2010; Sterner et al., 2020b). In the interest of collaboration and integration, it must also be able to expand to incorporate new organisms, new types of biological phenomenon, and so on.

Scientific understanding is of course ever-changing and expanding, and changing relationships between disciplines and fields requires constant revision of the scope and diversity of a bio-ontology’s structure and content. The continuing authority and trust afforded to the GO hinges critically on the mechanisms by which it responds to these changes (Leonelli, 2010; Leonelli et al. 2011).

3. Promises and Problems for Scientific Metaphysics

The previous section provided a brief sketch of the form and function of bio-ontologies: what they consist of, what purposes they serve, and how they are administered. It is due to these features of bio-ontologies that they may be a very promising avenue of research for scientific metaphysicians that has gone largely unexplored. (Not without exception; for example, Bechtel (2017) draws attention to bio-ontologies to lend scientific clarity to the longstanding problem of emergence.) Here I elaborate on the reasons for optimism about this project, which I prefaced in Section 1. However, each of these reasons comes with potential problems that complicate, and

may even undermine, the relationship between bio-ontologies and scientific metaphysics as a philosophical pursuit. I address each of these in turn.

The first promising feature is that, in aiming to capture biological knowledge in formal and explicit terms, the structure of bio-ontologies resembles Ontological theories in the philosophical sense. That is, they import concepts and principles developed in philosophical contexts to give explicit logical rigour to working scientific theories. It turns out that the kind of clarification of concepts in which metaphysicians have been engaged for some time is very important when outsourcing our reasoning to machines (Le Bihan & Barton, 2015): To avoid compounding errors in automated reasoning, one has to carefully work out inferential rules associated with those concepts and encode those rules into the system. For example, in the case of *continuants* (as opposed to *occurrents*, or processes) in OBO ontologies, the *is_a* relation must specify a time at which that relation obtains. Otherwise, this leads to errors: Every instance of *adult* was at one time an instance of *child*; it doesn't follow that *adult is_a child* (Smith et al. 2005).

What's more, the fact that bio-ontologies evolve in response to the changing state of scientific knowledge and the feedback of scientists who use them endows those ontologies with the empirical validity that scientific metaphysicians demand. At first glance, then, if one is interested in the Ontological structure of a particular biological domain, one could do no better than to look at the bio-ontologies that scientists investigating that domain choose to structure and manage their data, since they lay bare those Ontological commitments in full and explicit detail. For example, if one wishes to analyze a scientific theory from a structural realist perspective, a way to identify the theory's structure would be to look at the relational structure of the associated ontology or ontologies—that is, what relations are used to express knowledge of its domain, the logical rules associated with those relations, and so on. A scientific metaphysician interested in composition and part-whole relationships in biology, for example, might simply look at how the relation “is a part of” is used in the handling of scientific data in that context, what things are taken to stand in these relations with each other, the evidence demanded for claims about these part-whole relationships, and so on (Smith et al. 2005; Leonelli et al. 2011).

Second, the fact that bio-ontologies are deliberately designed for purposes of data integration makes them potentially highly relevant to metaphysical questions about pluralism and the unity of science. A common view in philosophy of science has been that science aims to establish a single, unified picture of the world—by the abandonment of inadequate theories, the reduction of higher-level or more specific theories to lower-level or more general ones, and so on (Oppenheim & Putnam, 1958). However, claims about the unity of science as a goal must deal with the fact of scientific pluralism—that equally respectable sciences can operate with very different, even apparently contradictory, theoretical pictures. Some take this diversity to be explained, even “fully justified” (Leonelli, 2010, p. 110), by the fact that a discipline's theory forms a complement with its own experimental methods, tacit knowledge, purposes, objects of study, and so on. In short, scientific pluralism as a philosophy generally interprets theoretical

diversity among the sciences as each adapting to its own practical context, and hence rejects the claim that science aims, or should aim, to unify under a single overarching theory. Rather than a sign of immaturity, the fact of plurality is a natural consequence of scientific progress in investigating a complex world (Dupré, 1993; Mitchell, 2003; Kellert et al., 2006; Waters, 2016).

An understanding of bio-ontologies and their interoperability may offer a new perspective on this issue. The increasing emphasis on integration of data from diverse origins challenges at least some versions of pluralism—those that claim that there is no scientific motivation to unify the diverse ontological pictures that scientists use. To participate in the wider community of data-centric science, it is no longer feasible for a particular community to simply hold to its own conceptual framework for its own idiosyncratic context without any consideration for how those concepts relate to those of other communities. As Laubichler et al. (2018) claim, the goal of integrating data produced in different disciplinary contexts requires us to consider clearly how their theoretical posits relate to each other—that is, to develop “a theory of theory integration”. Arguably, by abstracting from the idiosyncrasies of individual perspectives, the resulting integrated theory would be more objective (or at least intersubjective), since it would accommodate a wider range of scientific viewpoints. Importantly, this argument against strong pluralism is not just motivated by a priori intuitions about a unified reality, but by the concrete needs of science in practice. While this fact about bio-ontologies may not be taken to *entail* either a pluralist or monist metaphysical outlook (recall that I do not make any particular metaphysical claim here), it might at least put pressure on the hard-line pluralist by providing empirical observations about science that they must account for. In short, the process of integrating bio-ontologies forces us to engage with the idea that for all their differences, different scientific models, theories, and disciplines are in some sense dealing with the same thing (in the broadest possible sense of “thing”)².

Thirdly, and finally, another reason why bio-ontologies might be taken to be relevant to metaphysics is that prominent researchers consider the metaphysical interpretation of bio-ontologies to be critically important to their design. For them, the importance of metaphysics to bio-ontology design is not just about the clarification of concepts and their associated inferential rules. Instead, they see effective data management in science as depending critically on the adoption of a certain kind of metaphysical outlook. The OBO Foundry was co-founded by a metaphysician, Barry Smith, following a series of influential critiques of existing data systems used in biomedicine (e.g. Smith & Ceusters 2006) that led to the introduction of strict standards to govern those systems and their evolution. According to Smith (and others—see Munn & Smith 2013), the root cause of the design flaws in those earlier systems was precisely the designers’ rejection of the notion that data represent a reality independent of us and our concepts. To avoid those problems, they argue, one must be motivated by a particular kind of Realist outlook; namely, that the general kind terms of bio-ontologies refer not to mind-dependent concepts but to Aristotelian universals (see e.g. Smith, 2004). This discussion is likely to pique the interest of the scientific metaphysician: from their perspective, it would be a strong argument

in favour of some metaphysical theory if the adoption of that theory has been instrumental to scientific success.

Taken together, these three observations give us reason to see bio-ontologies as valuable sources of evidence in discussions about metaphysics. In general, they paint a picture in which bio-ontologies are, effectively, the empirical application of metaphysical ideas, and hence their uptake by scientists driven by empirical success counts as support for their metaphysical features. More strongly, these observations might even establish that bio-ontologies simply *are* metaphysical theories, or are at least underwritten in important ways by metaphysical commitments. If so, these theories gain critical naturalistic credibility from their demonstrable benefits to actual empirical work.

However, things are not so simple: each of the reasons above belie complications to the relationship between bio-ontologies and metaphysical Ontology. Each of these have to do with the practical, technological, social and epistemic factors that affect bio-ontology design in important ways.

Consider the first observation above—that bio-ontologies are clear formal conceptualizations of biological theories. The complication here is this: it's not clear that bio-ontology design is straightforwardly an exercise in expressing the consensus Ontology of the science, since the primary purpose for their design are the practical needs of efficient storage, sharing, integration, and so on. These practical needs may (and should, given those purposes) manifest in the design of the ontologies. For example, the fact that bio-ontologies must identify clear boundaries for their terms and strict rules for relations may simply follow from the nature of computational representation and the demand for reliable automated processing of large volumes of data. But the domains of biological and biomedical science may not neatly reflect the precise structures and rigorous characterizations that are to be found in the computational processes; in fact, the difficulties inherent in bio-ontology curation may be due to the fact that the world simply isn't fully understandable in terms of essential definitions, part-whole relationships, and so on (Dupré, 1993; Ladyman & Ross 2007). To put it one way, it may be that the behaviour of the biological world—being fuzzy, contingent, contextual, and exception-ridden—is more akin to natural language than to computational language; it's just that we can only use the latter *in silico*. In short, those features of bio-ontologies that make them potentially tractable to metaphysical analysis—their explicit formal structure and strict rules—may simply be an artifact of the technological role of bio-ontologies rather than a direct reflection of the Ontological structure of the domains they represent. Hence, we cannot straightforwardly conclude that the empirical domain of the science in question has some features from the fact that they are represented as having those features in bio-ontologies.

Second, consider the integrative role of bio-ontologies, which appear to make them pertinent to metaphysical questions about unification and reduction. The complication here is that the push towards unification and interoperability in bio-ontologies has a counteracting force; namely, the need to account for the diversity and dissensus that exists between the participating

communities. It is partly for this reason that not all scientific communities buy into calls for unification, and the associated formal strictures, of projects such as the OBO. Leonelli (2010) describes a looming threat of the perception of “terminological imperialism” in bio-ontologies: In order to maintain the trust of the users and the authority granted to the information the bio-ontologies provide, bio-ontologies must be responsive (and be seen to be responsive) to the individual needs of the participating communities rather than imposing standards that are alien to their particular practices and assumptions. It is partly this need that has given rise to the distinct profession of bio-ontology curators, whose job is to mediate between the requirements of the classification system and the individual scientific contexts that both produce and use the data. What’s more, it has been argued that unification in classification systems is not always necessary to facilitate integration efforts. As argued by Sterner et al. (2020a, b), integration doesn’t necessarily require forcing a consensus about classification where none exists; there are ways to pool data for integration purposes while allowing them to retain the idiosyncrasies of their origins, and to coordinate data representation in ways that honour dissensus in the definitions of terms.

These things complicate the relationship between integration in bio-ontologies and theoretical unification, since the claim that unification indicates metaphysical unity across domains is weakened if important domain-specific features are lost in the process. So in considering the Ontological implications of bio-ontologies, we need to consider not just what their interoperability enables but what it prohibits: as well as emphasising similarity, it may at the same time gloss over relevant differences in the actual phenomena³.

Finally, consider the third argument above—that many consider the metaphysical interpretation of bio-ontologies to be of demonstrable scientific importance. The complication here is that researchers involved in ontology design themselves disagree on the metaphysical commitments underlying design decisions. This is made clear by an exchange between Merrill (2010a,b) and Smith and Ceusters (2010) in the journal *Applied Ontology*. A central theme of that exchange is whether the design and administration of biological and biomedical ontologies is or should be based on a Realist methodology, or indeed on any metaphysical principles beyond the technical demands of logic, formal semantics, and philosophy of language. Merrill argues that the criticisms of and subsequent improvements to bio-ontologies attributable to Smith and Ceusters, while important, are also justifiable on non-realist grounds such as under a nominalist interpretation of kind terms. Further evidence of disagreement about the metaphysical underpinnings of ontology design is evident in discussions on the OBO message board⁴. The details of this exchange are too numerous and intricate to discuss here. The key point for present purposes is simply that practitioners disagree about whether and how we must interpret bio-ontologies in metaphysical terms to properly motivate and direct their design and administration. This challenges the idea that those metaphysical commitments are as essential to the concrete practical aspects of bio-ontology as some make them out to be, and may call into question whether there is any substantial metaphysical import to bio-ontologies at all.

These complications lead us to conclude that, at the very least, tying metaphysical discourse to the tangible features of scientific practice is a double-edged sword: While the concrete practical functions of scientific artifacts and practices such as those surrounding bio-ontologies are the source of their promise for scientific metaphysics, they are also the very thing complicating the goal of drawing metaphysical conclusions from them. This raises more general foundational questions affecting the prospects of scientific metaphysics: How can we defer to the epistemic and other features of science—those from which science draws its success—while retaining the metaphysical (and not simply practical or epistemological) import of the conclusions we reach; and, for that matter, without presuming to go “around” or “beyond” science in ways that are suspect from a naturalistic perspective? In the next section, I will point towards ways to tackle these questions in relation to bio-ontologies.

4. A Way Forward

How are we to understand the nature and extent of the connection between bio-ontologies as scientific artifacts and Ontology as a philosophical pursuit? Answers to this question will lie somewhere between two extremes. One extreme position may be that that scientific ontology design just *is* Ontology in the philosophical sense—perhaps the only Ontology worth doing—and that bio-ontologies are literally metaphysical theses about how the world hangs together. At the other extreme, it might be argued that the structure of bio-ontologies has nothing at all to do with the questions that concern metaphysicians—that despite the formal similarity, applied ontology as it relates to scientific data is entirely distinct from and irrelevant to Ontology in the philosophical sense. When we consider the complications created by the local, practical, epistemic, social and technological factors that affect how bio-ontologies are structured, we call the first extreme position into question, but potentially risk slipping into the second. The challenge, then, is to find a reasonable position between these two extremes. In this section, I suggest a way forward by proposing how to approach the metaphysical interpretation of bio-ontologies with the above three issues in mind.

Following the same order as above, consider first the point that bio-ontologies formally resemble Ontological theories (or frameworks for making Ontological claims). The complication I raised is that the design of bio-ontologies is adapted to their function as digital repositories of data and to structure automated reasoning, and so certain design features may simply be a result of those technical requirements rather than reflections of how the world being described is structured. At the very least, the existence of other explanations for certain structures in bio-ontologies, such as the need for reliable automated inferences, weakens the case for straightforward Ontological interpretations of those features. However, rather than being fatal to the idea of drawing metaphysical lessons from bio-ontologies, we can instead take the influence of these contextual practical factors as something we must consider before we jump to metaphysical conclusions. To be sure, bio-ontologies are in the first instance technical artifacts

designed by and for scientific communities with particular goals. However, it is critical to remember that many of those goals relate in important ways to successful investigation of the biological world, and so depend for their success on what the biological world is like. So while their design and successful application is certainly influenced by practical, social, technological, and other factors, they are presumably also affected in some way by the domain of inquiry as it is encountered in the process of scientific investigation. A task for scientific metaphysicians, then, is to tease out general principles (if indeed there are any to be found) about how these various influences, including the domain of inquiry, interact in the development of bio-ontological structures. Only then can they properly understand the relationship between those structures and the biological world they represent.

Similar lessons apply to the second issue—of data integration and its implications (or lack thereof) for questions about the unity of science. As I observed above, many consider the demands imposed by integration and interoperability to be in tension with the context-specific needs of individual communities. However, while there are situations in which features of bio-ontologies do not reflect a consensus (but rather a compromise, or perhaps even an executive decision that is not endorsed by all parties), we can take cases where there *is* consensus about ontologies and their interrelationships as being potentially metaphysically significant, even to a non-pluralist. Consensus is a matter of degree; the stronger the consensus, and more diverse the communities among which that consensus is reached, the more confident we can be that it is a matter of cross-contextual agreement about the world. What's more, disagreement isn't necessarily a sign of immature science on the way to becoming more settled and universalized; instead it may be that differences in the communities' respective domains demand that they be represented differently in data. Plurality, like unification, is potentially metaphysically significant, since pluralism is (for some) a metaphysical thesis.

This point might be understood as an argument about *robustness* in the sense of Wimsatt (2007): Biological phenomena are varyingly stable over time, across contexts, under different perspectives, and so on. When they are stable, and to the extent that they are, science will tend to track them with terms that are similarly robust—that is, more amenable to consensus between scientific contexts or perspectives. With that in mind, we can potentially interpret high degrees of consensus in bio-ontologies as evidence of relatively high robustness in the phenomena in question. In other words, consensus can be seen as a kind of *triangulation* of robust patterns in the biological world. (This evidence is defeasible, however, since categories may be stable not because they capture a robust feature of the world but because they are developmentally entrenched⁵.) More generally, this kind of reasoning points the way to a critical point of contact between metaphysics and epistemology that scientific metaphysicians require. That is, when we understand epistemic features (such as consensus) as functionally coupled to certain features of the world (such as robustness), this permits inferences between the two.

Finally, consider the issue of disagreements about the metaphysical interpretation of bio-ontology, as represented in the exchange between Smith and Ceusters and Merrill, for

example. I will not attempt to settle this dispute here. Instead, the critical point for the present discussion is this: For all their differences, participants in this debate seem to agree that it is only important to the extent that it has implications for scientific methodology. That is, philosophical discussions about things like universals are important to science when, and only when, they make a difference to how scientific institutions design, administer, regulate and improve scientific ontologies, and when those differences affect the success of the scientific enterprise. Smith and Ceusters, for example, claim that understanding terms as corresponding to things in reality places a critical constraint on what terms are admissible, which is necessary to facilitate integration. If there were another, non-realist way to motivate the same constraint, as Merrill claims, the discussion has at that point passed beyond the horizon of empirical relevance, and into what Amanda Bryant (2020) calls “free range metaphysics”. On the other hand, if those constraints on the admissibility of terms are shown to be too strict for concrete practical purposes—for instance, if they lead ontologists to disallow terms such as “non-smoker” or “absent arm” that are used by practitioners⁶—then one must either relinquish those realist commitments or, if one prefers, find a realist-friendly way to interpret those terms in the formal ontology. Either way, the result is by hypothesis the same: what ultimately matters is what is done, not the “ism” we attach to what is done.

This point of agreement is important for understanding how formal scientific ontologies and their integration matter to scientific metaphysics in the present sense. For scientific metaphysics to be properly empirical (hence earning its *scientific* epithet), it should be attentive to the horizon within which metaphysical principles have concrete practical and empirical implications, and hence can potentially make a difference to success and failure in scientific practice. For only when they do make a difference to practice, and hence potentially to success, can they be said to be justifiable on scientific grounds. Precisely where this horizon of practical relevance lies is evidently undecided even among practitioners. Whether there is any value in general to super-empirical disputes that lie beyond the horizon of the concrete needs and purposes of science, this is irrelevant by definition to scientific metaphysics as an empirically-constrained project. In any case, discussions about the proper way to design and regulate bio-ontologies are of direct importance to scientific metaphysicians, for they are effectively discussions about which metaphysical disputes are scientific at all.

5. Discussion: What is Scientific Metaphysics?

As mentioned, the project of scientific metaphysics is founded on the principle that, since science is so successful at predicting, explaining and manipulating the world, its success-making features should tell us something about what the world is like. Yet this leaves open exactly how this project is to be carried out—which aspects of science are metaphysically relevant, and how observations about those aspects are to be translated into metaphysical claims. In this section, I

draw from the above discussion about the metaphysical interpretation of bio-ontologies to address broader questions about scientific metaphysics in general.

Scientific metaphysicians are agreed that metaphysics should draw on successful science. However, when we investigate concrete cases of success-making scientific phenomena, such as bio-ontologies, drawing metaphysical inferences from those features looks far from straightforward. As we've seen, bio-ontologies—indeed, working scientific theories in general—are embedded in a complex web of practical goals, social interactions, technological needs and constraints, and trade-offs between all of these. Because of this, we cannot straightforwardly infer from the fact that a bio-ontology defines its entities and relations in a certain way, for example, to the conclusion that the domain's entities and relations really have those features (with “really” interpreted in some strong metaphysical sense). However, as I argued in the previous section, this doesn't mean that scientific metaphysics is a lost cause. What it does mean is that scientific metaphysics cannot ignore those factors, and must instead closely study how they contribute to the formation and evolution of science's ontologies—both its formal, digital ontologies and the relatively informal ones used in journal articles or lab discussions. Only then can we bring into view what those ontologies say about the world that those sciences are investigating.

In short, the complexities apparent in cases such as that of bio-ontologies call for a scientific metaphysics that explicitly studies and incorporates scientific practice. Some philosophers have taken this approach, more or less implicitly, for some years now. However, there have recently been focused attempts to galvanize researchers in this line of inquiry (Wimsatt 2007; Waters 2016; Ereshefsky and Reydon, forthcoming) The present work also aims to embody this mission. With this in mind, we can address three questions—those that any scientific metaphysics should address—with the lessons of the case of bio-ontologies in mind. The answers I propose all invoke the importance of attention to scientific practice.

There are, as I see it, a number of distinct questions about scientific metaphysics to which any version of it should provide an answer. First, what is *scientific* about one's scientific metaphysics? That is, in what way does it draw on science to inform or otherwise constrain its conclusions? Second, what is *metaphysical* about it? That is, in what sense is it a distinct pursuit from both science and “mere” epistemology? In distinguishing itself from science and epistemology, it must not (given the broadly naturalistic motivations of scientific metaphysics) presume to be going beyond or around science in ways that are problematic from a naturalistic perspective. Finally, what is the ultimate aim or purpose of this scientific metaphysics? What intellectual or other kind of goals does it serve? My above discussion of the metaphysical significance of bio-ontologies yields more general lessons for how we might understand the nature, methods, and aims of scientific metaphysics, which I discuss here.

Consider the first question—of what observations about science might inform metaphysical questions, and why. Bio-ontologies, as I've shown, are a strong candidate for a source of fodder for metaphysical discussion because their design features explicitly incorporate

concepts developed in metaphysics and other branches of philosophy. While bio-ontologies are certainly influenced by their practical contexts in a variety of important ways, this does not preclude their being used by scientific metaphysicians to inform their own philosophical discussions about Ontology. If we understand metaphysics as the search for the best conceptualization(s) of the world we live in (whether by “best” we mean the truest, or something else), then we would do well as scientific metaphysicians to look at how scientists see fit to conceptualize it. Practical complexities notwithstanding, bio-ontologies do represent varying degrees of consensus between scientific communities about how to conceptualize their domains for the purposes of predicting, explaining, and manipulating them. There are, as we have seen, other factors affecting the structure of bio-ontologies besides simply “what the world is really like”. Rather than ignoring this, a properly scientific metaphysics should be acutely aware of those influences in order to control for them—to determine the extent to which the structure of a bio-ontology says something about the structure of the domain it is used to represent.

On the second question—of what is metaphysical, and not “merely” scientific or epistemological about this project—the answer can only be a complex and nuanced one. As I take it, the very possibility of scientific metaphysics implies the rejection of sharp divides between the epistemological and the metaphysical and between the metaphysical and the scientific, since it is precisely the epistemic successes of science that are taken to underwrite its metaphysical significance. A properly scientific metaphysics goes *through* science, not around it. Yet this leaves open exactly how scientific metaphysics stands apart from those other pursuits.

In this discussion, we can see a relatively clear picture of where metaphysics enters the realm of science, and vice versa, while remaining distinct. As I’ve mentioned, a longstanding view of metaphysics is that it is independent of scientific findings because it is the very conceptual framework within which empirical observations can be understood. In a way, that is what bio-ontologies are: they are “prior” to empirical findings in that they provide the very conceptual structure within which those empirical findings are represented. However, this does not make bio-ontology design a purely a priori exercise: As science develops, communities of scientists adjust those frameworks to better suit their evolving empirical needs. In that sense, then, they are not prior to empirical data but an attempt to accommodate and organize them that evolves as those data accumulate. In short, the study of bio-ontologies from a metaphysical perspective may be better suited to a relativized a priori account of the role of metaphysics in science: In brief, scientific metaphysics is concerned with the conceptual or logical frameworks within which science structures and organizes its data and, critically, how those frameworks change in parallel with the scientific enterprises they support. A key challenge for scientific metaphysics is to understand the mechanisms by which those changes occur.

Aside from questions about the relationship between metaphysics and science, we also need an account of the relationship between metaphysics and epistemology—one that establishes our pursuit as a case for the former. I’ve argued that, in the process of analyzing bio-ontologies for metaphysical purposes, one must carefully consider the various factors that influence their

structure, which include the epistemic, practical, social and technological contexts they are designed for. However, we can suppose that the best way to classify a domain in the form of a bio-ontology is at least partly determined by what that domain is like, *as well as* those other factors. This approach generalizes beyond bio-ontologies: Consider Alan Love's claim that typology in developmental biology is not a metaphysical claim but instead "a form of *thinking* or reasoning, as conceptual behavior" (Love, 2009, p. 53). This may well be true, but this doesn't prevent us from supposing that the value of typological thinking in developmental biology has something to do with the world: If the workings of biological development had been different, it might not have been so useful to scientists to think about it typologically. In short, while we should certainly understand scientific classifications such as those of bio-ontologies in epistemic terms—as modes of thinking—we can also take those modes of thinking to be more or less well-suited to the investigation of certain types of phenomena. To say that an object "makes a good hammer", while essentially tied to intended use, also says something about that object: that it is heavy, that it isn't brittle, and so on. In the same way, to say that a biological phenomenon can be usefully studied with this or that epistemic strategy is to say something about that phenomenon.

This discussion leads to an important principle underlying the emerging project of incorporating biological practice into discussions of scientific metaphysics. Typically, philosophical work that emphasises scientific practice either brackets or actively distances itself from metaphysical questions, focusing instead on its epistemic, social, technological, and other features. Yet it is beginning to be realized that while these features are indeed the source of science's success, this success is not fully accounted for until we consider the role of the world in making scientific practices successful (Waters 2016; Ereshefsky and Reydon, forthcoming). Scientific interactions with the world are *functional* interactions, and that functionality involves not just "internal" features of science such as its theories and methods, but also the external environment with which they interact. This key role for the world in determining the success of those interactions is what justifies these projects as metaphysical. A metaphysics based on scientific practice is tasked with understanding those interactions.

To put this another way, we can understand scientific epistemology and scientific metaphysics as effectively studying the same thing—scientific practice—but with different emphasis. Both are concerned with the nature and extent of scientific consensus about how best to conceptualize the world in order to successfully interact with it. However, each focuses on different aspects of what makes one way better than another: While epistemology typically emphasises the social and epistemic forces affecting the development of biological knowledge, metaphysics considers the extent to which that development is affected and constrained by the biological world being investigated. Any attempt to do one without the other, I argue, is philosophically suspect: we cannot make sense of the idea that we should aim to express what the world is like *rather than* what we believe about what the world is like. Claiming to do so leads us to a pragmatic contradiction (van Fraassen, 2008, Ch. 11).

Finally, this brings us to the question of the aims of scientific metaphysics: what intellectual purpose is served by the project of teasing apart the different influences on epistemic artifacts in science such as bio-ontologies? Ladyman and Ross (2007) argue that the proper role of scientific metaphysics is to unify the sciences by relating them all to the posits of fundamental physics, thus placing unilateral constraints on the admissible claims of the special sciences. While this is certainly one legitimate role, there is perhaps a more general view we can take. First, there are ways to illuminate the relationships between scientific disciplines besides relating them all to fundamental physics; in the OBO, for example, the unifying theory is that of the Basic Formal Ontology, which provides the conceptual building blocks for the more provincial ontologies. With that in mind, the role of philosophers and philosophical thinking in that project is unsurprising—it can be understood in terms of what some see as the broader aim of philosophy, which is to provide the connective tissue between other disciplines (Hoffmann, Schmidt & Nersessian 2013). As we've seen, this requires understanding how each discipline's theoretical elements are adapted to their particular contexts, methods and aims, so that those influences can be abstracted in order to establish a shared conceptual basis that all parties can agree on.

Second, there is value not just in clarifying the connections between scientific disciplines, but also between science and other spheres of human interest. It's easy to see how one might be interested in what science has to say about what are often considered metaphysical matters—about individuality, causality, sex, gender, race, parthood, free will, mind and matter, and so on—for purposes that are not themselves scientific. While science is of course deeply interconnected with other spheres of human activity and the world at large, we can understand it to have aims that are in some sense intrinsic to it—aims relating to prediction and control, explanation, revealing new avenues of investigation, and so on. In fact, this difference in purpose is what complexifies the relationship between how science treats those matters and how humanity more generally should think of them—why scientific metaphysics is not a straightforward matter of “reading off” its conclusions from the sciences. As Haslanger (2016) argues, for example, the standards by which the biological sciences sort human beings into sex categories may not be well-suited to other human institutions or to our intrapersonal self-understanding, simply because the aims for their categorizations are different. Nevertheless, these different spheres are in an important sense overlapping, and so we can consider one task of scientific metaphysics as aiming to explore and clarify the extent and implications of this overlap. Just like the project of mediating between scientific disciplines, a necessary part of this wider project will be to understand how scientific ontologies are affected by the sciences' particular interests and purposes, and hence how other spheres of human inquiry might employ ontologies more suited to their own. This, in short, is precisely the project of scientific metaphysics I have outlined here.

6. Conclusion

In this paper, I outlined some general features of the form and function of bio-ontologies used in modern data-driven biological and biomedical science. At first glance, they may appear to be straightforwardly interpreted as Ontological frameworks used by the sciences, and therefore ideal sources of evidence for scientific metaphysicians interested in arbitrating between metaphysical claims by appeal to successful science. However, there are some important factors complicating that interpretation, which all have to do with the fact that bio-ontologies are in the first instance artifacts designed for concrete practical purposes, some of which weaken the connection between them and metaphysics as traditionally understood. This is not fatal to the idea that bio-ontologies can be analyzed in the service of metaphysics, however. In fact, the idea that bio-ontologies are influenced by their practical and epistemic context makes them quite amenable to analysis from a scientific metaphysics perspective. Since scientific metaphysics takes epistemic success to be the source of science's metaphysical significance, a rigorous scientific metaphysics should simply aim to understand how successful scientific practices, and their associated artifacts such as bio-ontologies, are influenced both by their local epistemic contexts and by the domain they are designed to investigate.

In this work I have explicitly avoided committing to a particular metaphysical theory or approach in proposing how bio-ontologies might be studied from a scientific metaphysics perspective. Which theory or approach is best is left to the outcome of future work—work that I hope to have convinced the reader is worthwhile. The essential role of attention to scientific practice is the only firm recommendation that I offer to this ongoing project.

Notes

¹ A related field is what is often called *metaphysics of science*, which (as I understand it) aims to establish the metaphysical underpinnings of concepts that are employed in scientific theories, such as *causation* and *law*. I take this to be a somewhat different endeavour than scientific metaphysics, though they may overlap.

² Brigandt (2010) argues that unification is often localized to particular research aims rather than reflecting a global move towards the ideal of a unified science. However, the idea that theoretical unification is both possible and desirable in local contexts may be cause for optimism about the possibility of a more far-reaching unification project, even if scientists themselves have more local aims.

³ Of course (as pointed out by an anonymous reviewer), adherents to a thoroughly pluralist metaphysics are likely to be untroubled by this problem, since they do not see unification as a goal of science or a requirement for its metaphysical relevance in the first place. This point

reminds us that disagreements about metaphysics are difficult to separate from disagreements about what constitutes a metaphysics at all.

⁴ <http://tinyurl.com/34lacvy>

⁵ Thanks to an anonymous referee for pointing this out.

⁶ For interesting informal discussion on these issues in the OBO discussion group, see <http://tinyurl.com/34lacvy>.

References

- Bechtel, William (2017) Using the hierarchy of biological ontologies to identify mechanisms in flat networks. *Biology & Philosophy* 32: 627–649
- Bechtel, William (2020) Data Journeys Beyond Databases in Systems Biology: Cytoscape and NDEx. In S. Leonelli and N. Tempini (Eds.) *Data Journeys in the Sciences* (pp. 121–43). Springer.
- Bennett, Karen (2016) There Is No Special Problem with Metaphysics. *Philosophical Studies*, 173 (1): 21–37.
- Berg, Ellen L. (2014) Systems Biology in Drug Discovery and Development. *Drug Discovery Today*, 19: 113–25.
- Brigandt, Ingo. (2013) Integration in Biology: Philosophical Perspectives on the Dynamics of Interdisciplinarity. *Studies in History and Philosophy of Science Part C :Studies in History and Philosophy of Biological and Biomedical Sciences*, 44, 461–65.
- Bryant, Amanda. (2020) Keep the Chickens Cooped: The Epistemic Inadequacy of Free Range Metaphysics. *Synthese*, 197, 1867–1887.
- Chakravartty, Anjan. (2013) On the Prospects of Naturalized Metaphysics. In D. Ross, J. Ladyman, and H. Kincaid (Eds.) *Scientific Metaphysics* (pp. 27–50). Oxford University Press.
- Daston, Lorraine (2004) Type Specimens and Scientific Memory. *Critical Inquiry*, 31, 153–82.
- Dupré, John. (1993) *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Cambridge, MA: Harvard University Press.
- Ereshefsky, M. and Reydon, T.A.C. (forthcoming) The Grounded Functionality Account of Natural Kinds. In Bausman, W., Baxter, J., Lean, O., Love, A. & Waters, C.K. (Eds.) *From Biological Practice to Scientific Metaphysics*. Minnesota Studies for Philosophy of Science. University of Minnesota Press.
- Friedman, M. (2001) *Dynamics of Reason*, Stanford: CSLI Publications.
- Haslanger, Sally (2016) Theorizing with a Purpose: The Many Kinds of Sex. In C. Kendig (Ed.), *Natural Kinds and Classification in Scientific Practice* (pp. 129–44). Routledge.
- Hirschman, L, Burns GAPC, Krallinger M, et al. (2012) Text mining for the biocuration workflow. *Database*, bas020.

- Hoehndorf, Robert, Paul N Schofield, and Georgios V Gkoutos. (2015) The Role of Ontologies in Biological and Biomedical Research: A Functional Perspective. *Briefings in Bioinformatics* 16, 1069–80.
- Hoffmann, Michael HG, Jan C Schmidt, and Nancy J Nersessian (2013) Philosophy of and as Interdisciplinarity. *Synthese*, 190, 1857–64..
- Kellert, Stephen, Helen Longino, and C Kenneth Waters. (2006) The Pluralist Stance. In S. Kellert, H. Longino and C. K. Waters (Eds.) *Scientific Pluralism* (pp. vii–xxix) University of Minnesota Press.
- Kincaid, Harold (2013) Introduction: Pursuing a Naturalist Metaphysics. In D. Ross, J. Ladyman, and H. Kincaid (Eds.) *Scientific Metaphysics* (pp. 1–26). Oxford University Press.
- Ladyman, James, and Don Ross (2007) *Every Thing Must Go*. Oxford University Press.
- Laubichler, Manfred D., Sonja J. Prohaska, and Peter F. Stadler (2018) Toward a mechanistic explanation of phenotypic evolution: The need for a theory of theory integration. *The Journal of Experimental Zoology Part B: Molecular and Developmental Evolution*, 330, 5–14.
- Le Bihan, Baptiste, and Adrien Barton (2018) Analytic Metaphysics Versus Naturalized Metaphysics: The Relevance of Applied Ontology. *Erkenntnis*. doi.org/10.1007/s10670-018-0091-8
- Leonelli, Sabina (2010) Documenting the Emergence of Bio-Ontologies: Or, Why Researching Bioinformatics Requires HPSSB. *History and Philosophy of the Life Sciences*, 32, 105–25.
- Leonelli, Sabina (2015) What Counts as Scientific Data? A Relational Framework. *Philosophy of Science* 82, 810–21.
- Leonelli, Sabina (2016) *Data-Centric Biology: A Philosophical Study*. University of Chicago Press.
- Leonelli, Sabina, Alexander D. Diehl, Karen R. Christie, Midori A. Harris, and Jane Lomax (2011) “How the Gene Ontology Evolves.” *BMC Bioinformatics* 12, 325.
- Love, Alan C. (2009) Typology Reconfigured: From the Metaphysics of Essentialism to the Epistemology of Representation. *Acta Biotheoretica*, 57, 51–75.
- Lowe, E. J. (2002) *Survey of Metaphysics*. Oxford: Oxford University Press.
- Merrill, Gary H. (2010a) Ontological Realism: Methodology or Misdirection? *Applied Ontology*, 5, 79–108.
- Merrill, Gary H. (2010b) Realism and Reference Ontologies: Considerations, Reflections and Problems. *Applied Ontology*, 5, 189–221.
- Mitchell, Sandra D. (2003) *Biological Complexity and Integrative Pluralism*. Cambridge University Press.
- Munn, Katherine, and Barry Smith (Eds.) (2008) *Applied Ontology: An Introduction*. De Gruyter.
- Oppenheim, P. and H. Putnam (1958) The unity of science as a working hypothesis. In H. Feigl et al. (Eds.), *Minnesota Studies in the Philosophy of Science*, vol. 2, Minneapolis: Minnesota University Press.

- Paul, L. A. (2012). Metaphysics as modeling: The handmaiden's tale. *Philosophical Studies*, 160, 1–29.
- Searls, David B. (2005) Data Integration: Challenges for Drug Discovery. *Nature Reviews Drug Discovery*, 4, 45–58.
- Smith, Barry (2003) “Ontology.” In L. Floridi (Ed.) *Blackwell Guide to the Philosophy of Computing and Information* (pp. 155–66). Blackwell Publishing.
- Smith, Barry (2004) Beyond Concepts: Ontology as Reality Representation. In A Varzi and L. Vieu (Eds.) *Proceedings of FOIS 2004, Turin*.
- Smith, Barry, and Werner Ceusters (2006) HL7 RIM: An Incoherent Standard. In A. Hausman, R. Haux, J. van der Lei, E. de Clercq, and F. H. Roger France (Eds.), *Ubiquity: Technologies for Better Health in Aging Societies*. IOS Press.
- Smith, Barry, and Werner Ceusters (2010) Ontological Realism: A Methodology for Coordinated Evolution of Scientific Ontologies. *Applied Ontology*, 5, 139–88.
- Smith, Barry, Werner Ceusters, Bert Klagges, Jacob Köhler, Anand Kumar, Jane Lomax, Chris Mungall, Fabian Neuhaus, Alan L Rector, and Cornelius Rosse (2005) Relations in Biomedical Ontologies. *Genome Biology*, 6, R46.
- Smith, Barry, Michael Ashburner, Cornelius Rosse, Jonathan Bard, William Bug, Werner Ceusters, Louis J. Goldberg, et al. (2007) The OBO Foundry: Coordinated Evolution of Ontologies to Support Biomedical Data Integration. *Nature Biotechnology*, 25, 1251–55.
- Sterner, Beckett, Steve Elliott, Edward Gilbert, and Niko Franz (2020a) “Data Integration without Unification.” <http://philsci-archive.pitt.edu/19253/>
- Sterner, Beckett, Joeri Witteveen, and Nico Franz (2020b) Coordinating Dissent as an Alternative to Consensus Classification: Insights from Systematics for Bio-Ontologies. *History and Philosophy of the Life Sciences*, 42, 1–25.
- van Fraassen, Bas C. (2008) *Scientific Representation*. Oxford: Oxford University Press.
- Waters, C Kenneth (2016) No General Structure. In M. Slater and Z. Yudell (Eds.) *Metaphysics in Philosophy of Science*. Oxford: Oxford University Press.
- Wimsatt, William C. (2007) *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Cambridge and London: Harvard University Press.