

Published in final edited form as:

Appl Ontol. 2010 November 15; 5(3-4): 139–188. doi:10.3233/AO-2010-0079.

Ontological realism: A methodology for coordinated evolution of scientific ontologies

Barry Smith* and Werner Ceusters

New York State Center of Excellence in Bioinformatics and Life Sciences, University at Buffalo, Buffalo, NY, USA

Abstract

Since 2002 we have been testing and refining a methodology for ontology development that is now being used by multiple groups of researchers in different life science domains. Gary Merrill, in a recent paper in this journal, describes some of the reasons why this methodology has been found attractive by researchers in the biological and biomedical sciences. At the same time he assails the methodology on philosophical grounds, focusing specifically on our recommendation that ontologies developed for scientific purposes should be constructed in such a way that their terms are seen as referring to what we call *universals* or *types* in reality. As we show, Merrill's critique is of little relevance to the success of our realist project, since it not only reveals no actual errors in our work but also criticizes views on universals that we do not in fact hold. However, it nonetheless provides us with a valuable opportunity to clarify the realist methodology, and to show how some of its principles are being applied, especially within the framework of the OBO (Open Biomedical Ontologies) Foundry initiative.

1. The methodology of ontological realism

1.1. The goal of ontology development

Ontologies are created to serve multiple goals, including support for more effective retrieval of data and for different sorts of reasoning. Here we focus on ontologies created to foster consistency in the ways scientific results are described for purposes of more effective integration of scientific data – ontologies, therefore, that serve strategies to counteract the many tendencies leading to *ad hoc* and non-interoperable coding of data, and thus to the formation of data silos.

Unfortunately, the very success of such strategies has led to the creation of ever new ontologies, and thus has resurrected the very silo problems which ontologies were designed to counteract. To this end, it is of obvious advantage if we can find a way to *minimize the number of ontologies* that are being constructed and at the same time *maximize their mutual consistency*. These ends can be achieved, however, only if we can persuade ontology developers to accept certain common *constraints on how they build their ontologies* and if we can find a way to do this that does not endanger the flexibility that is needed to keep pace with scientific advance.

The realist methodology is based on the idea that the most effective way to ensure mutual consistency of ontologies over time and to ensure that ontologies are maintained in such a

way as to keep pace with advances in empirical research is to view ontologies as representations of the reality that is described by science. This is the *fundamental principle* of ontological realism.

To ensure the wide use of each single ontology that is the presupposition for its serving the purposes of data integration we thus advocate the creation of a small and highly constrained system of *reference ontologies* designed to embody the representational content of settled science (Rosse & Mejino, 2003). In areas where flexibility is needed, for example where research is still exploratory and results provisional, we advocate the creation of *application ontologies* that are built as far as possible as extensions of corresponding reference ontologies.

In Section 1 of this essay we shall seek to describe the methodology in such a way as to bring out its practical significance.¹ In the remaining sections we address some of the philosophical issues raised by Merrill's critique.

1.2. Types, instances and resemblance

Different forms of realism are distinguished by philosophers, of which the most important for our purposes here are:

Scientific realism =_{def.} the doctrine according to which scientific theories are (broadly) true of reality.

Metaphysical realism =_{def.} the doctrine according to which universals or types exist in reality.

Merrill (2010, p. 85), quite correctly, sees elements of both of the above in the etiology of our thinking on ontology development. He himself embraces what he calls an 'anti-realist' position which consists in the denial of metaphysical realism as defined above, and which we can accordingly define as follows:

Anti-realism =_{def.} the doctrine according to which there are no universals or types in reality, but only individuals or particulars.

Two forms of anti-realism can then be distinguished:

Nominalism =_{def.} a variety of anti-realism consisting in a doctrine to the effect that entities labeled by the same term – for example, this bonobo and that bonobo – have nothing in common but their name.

Conceptualism =_{def.} a variety of anti-realism consisting in a doctrine to the effect that entities conceptualized in the same way have nothing in common but the fact that they are so conceptualized.

Disputes between the realist and anti-realist camps have raged for some thousands of years. Anti-realists object to the metaphysical realist position because they find appeals to entities such as universals or types unscientific. Metaphysical realists object to anti-realism (in either version) because they see it as involving its own tacit appeal to universals in reality (either in the realm of words and utterances, or in the realm of cognition).

Since 2002 we have been attempting to move beyond such disputes by developing a *methodology*, which we call 'ontological realism', that will capture what we believe to be a kernel of practical significance in these debates by addressing the question what it is to

¹The principles propounded in what follows are derived from our own practice in ontology development, and go beyond the principles thus far adopted by the OBO Foundry, which are documented here: <http://obofoundry.org/crit.shtml>.

which the terms used in ontologies should be seen as referring. Because ontological realism is a methodology, and not a doctrine, it stands in no logical relation to any of the metaphysical doctrines specified above. Certainly it takes over the terminology of ‘types’, ‘universals’, ‘instantiation’ from the metaphysical realist literature; but it does not stand or fall according to whether universals or types do or do not exist in some metaphysical sense, and our goal will be to provide a specification of our methodology which will allow even anti-realists to recognize its benefits.

1.3. The methodology

The methodology can be summarized as follows. Ontologists, when building ontologies, should conceive the world as including entities of two sorts – called ‘particulars’ (or ‘instances’) and ‘types’ (or ‘universals’), respectively. Particulars, according to this doctrine, are the sorts of things that can be described on the basis of observations performed for example in the lab or clinic. Types or universals – we shall always use these terms synonymously in what follows – are to be understood as counterparts in reality of (some of) the general terms used in the formulation of scientific theories. Particulars are concrete individual entities (entities that exist in space and time and that exist only once); types or universals are to be understood as *repeatable*. This means that, for each given type, we can in principle discover of indefinitely many particulars that they are its instances. (We shall return to address in more detail the relation between universals and repeatables below.)

The particulars in reality can be partitioned into groups on the basis of multiple similarity relations which obtain between them, and the process of recognizing such collections of similars is essential to all forms of cognition. Sometimes the process yields *classifications*, which is to say partitions of reality based on hierarchies organized in terms of greater and lesser generality.

Multiple more or less *ad hoc* classifications have been created in the course of time, and human beings have the ability to cope with the resultant mismatches. Computers, however, are much less tolerant of classificatory inconsistency, and this can cause problems when computers are put to work in managing large and heterogeneous bodies of data. We can distinguish two kinds of responses to these problems, of which the first, sometimes called the *bottom-up* approach, sees the solution in terms of mappings between the existing, mismatched classifications. The second, *top-down* approach, sees the solution in terms of strategies to constrain the classifications created and used by different groups, in the direction of greater consistency. The realist methodology that we advocate falls within this second camp, and its strategy of prospective standardization in some ways parallels the earlier effort to coordinate the expression of measurement results by creating a single international system of units.

Whether scientists themselves see the general terms they use as referring to types (or universals or natural kinds or like entities) is not relevant to the success of our methodology. All that is important is that scientists use general terms in attempting to describe repeatable features of reality. That they do this is not – Merrill’s odd imputations to us of views to the contrary notwithstanding – because they have been taught (or should be taught) special metaphysico-semantical doctrines concerning reference and meaning. Rather it is for a variety of practical reasons – for example, because, when scientists formulate particular and general assertions, then they want other scientists to be able to verify or falsify them in experiments. For this they must be in a position to describe repeatable features of reality in a way that allows these other scientists to recreate them.

We do not deny that there are many distinct philosophical approaches to the understanding of the scientific use of general terms and of what it is in reality towards which such terms

are directed. For practical purposes, however, we believe that these philosophical matters are of secondary importance. This is because even the metaphysical anti-realist can, we believe, view all putative references to types or universals – including the many such putative references in what follows – as mere *façons de parler* about other, more commonplace entities – such as *scientists' beliefs* or *linguistic usage* – and still gain full practical advantage from our methodology.²

We take as our starting point a distinction between two sorts of descriptions, which we believe pervades the whole of science. It is seen most simply in the contrast between, for example,

(A) AIDS is spreading very rapidly through Asia,

and

(B) AIDS is caused by the HIV virus,

in which the string 'AIDS' can be understood as referring to a particular collection (in (A)), and to a type (in (B)).

Scientists are constantly drawing on this distinction as they move back and forth between *descriptions of experiments* on the one hand, where they are dealing with carefully demarcated collections of particulars (for example, populations of study organisms), and the *formulation of results in theories* on the other hand, where they can be seen as dealing with corresponding types.

The distinction between collections and types is used by scientists themselves to monitor progress in discovering the structure of reality. It was a scientific advance when members of the collection of human beings distinguished by the possession of the phenotypic feature of having a mongoloid face were found to be associated with instances of what the realist would call the disorder universal *trisomy 21*. Similarly, the scientific debate over whether there exists something that is properly to be called a 'race' can be formulated in terms of whether 'race' should be understood as denoting a universal or a mere collection.

Discovering universals – for example, discovering that there is a type of disease called 'AIDS', or a type of particle called 'Higgs boson' – is a scientific achievement. Discovering that terms purportedly referring to universals (like 'diabetes') do not do so (or do so only ambiguously, as between 'diabetes mellitus' and 'diabetes insipidus') is a scientific achievement of a different kind. Yet another kind of achievement consists in discovering general truths about universals – for example, discovering that infection with the influenza virus causes the same type of disease throughout the world, even in spite of the many different manifestations and culturally contingent descriptions with which it is associated.

1.4. What are types or universals?

The difference between collections of particulars on the one hand and types or universals on the other is related to what is commonly referred to in some logical circles as that between *classes in extension* – roughly, sets of individuals in the familiar sense – and *classes in intension* – the latter sometimes (on one of the several understandings of this word) called 'concepts'. The problem with the approach in terms of extensions and intensions is that it

²Very crudely, the anti-realist might view a sentence of the form

(1) scientist *X* believes that instances of a given type *Y* exist in reality

as meaning something like

(2) scientist *X* believes that it is appropriate to use the general term '*Y*' in making positive assertions about reality.

suggests that there is a closer concordance between the two sorts of entities than is in fact the case. Both extensions and intensions, on standard views, can be combined in arbitrary ways in Boolean combinations. Thus if F is an extension (set) and G is an extension (set), then there are further extensions F & G, F or G, non-F, non-G, non-F & non-G and so on. And similarly: if F is an intension, for example, the concept *nausea*, and G is an intension, for example, the concept *vomiting*, then there are further concepts *nausea and vomiting*, *nausea or vomiting*, *non-nausea*, *non-vomiting*, *non-nausea and non-vomiting* and so on. Concepts, in other words, can be combined logically to produce other concepts. The uncontrolled combination of concepts in this manner is in our eyes one reason for the failure, thus far, of terminology artifacts created in accordance with what we shall recognize below as the ‘concept orientation’. This is because, from the potentially infinite number of concept combinations that can be formed from any given starting point, some selection must be made. And as different individuals and groups make their selections in more or less deliberate and more or less *ad hoc* ways, the realization of the goal of ontology-based integration becomes ever more remote.

Something similar holds, of course, for *collections*. Thus, for example, if there is a collection of people suffering from nausea in a given hospital at a given time, and a collection of people suffering from vomiting in the same hospital at the same time, then there are *ipso facto* further collections formed, for example, by the union and the intersection of these two collections.

As concerns types, however, things are different. If we know (or better: think we know, modulo the current state of science) that certain types exist, then on our view – which corresponds to the ways scientists themselves often use words like ‘type’ or ‘kind’ in devising terminologies to describe their results – the rules of Boolean algebra give us no sanction at all to infer that certain other types exist also. The question thus arises as to which collections *do* correspond to types or universals in the sense that we can formulate for them definitions of the following sort:

(C) Collection of X’s =_{def.} collection of particulars of type X.

This question is, unfortunately, not answerable with any simple recipe. It is in this respect comparable to the question: how do we establish whether a given scientific assertion is true? It would of course be nice to have a decision procedure for determining which terms should be recognized as designating types for any given discipline – ideally one which could be programmed into a computer. In fact, however, the set of candidate terms designating types is a matter that is decided, for each science, by the scientists themselves, in an on-going process of terminology evolution through which those terms come to be selected for that are fit to serve in successive formulations of the corresponding scientific theory. The work of the ontologist, as we see it, is in large part one of transforming the results of this process – which are standardly informal, unreflected, subject to redundancies, ambiguities and to constant revision – into the sorts of systematic representations that are needed to support integration.

Each scientific theory as it exists at any given stage will likely be marked by (as yet unidentified) terminologically relevant errors, and these errors will accordingly be carried over into the corresponding ontology. Hence, we cannot embrace any ‘representational assumption’ according to which there is a one-one correspondence either between scientific general terms, or between terms in reference ontologies, and types or universals in reality. Rather, the realist methodology is one according to which the developers of a reference ontology should *assume* for heuristic purposes that the terms in the ontology they are developing refer to such types, knowing full well that this assumption may be false for any given term. Ceusters (2009) shows how on this basis we can use the analysis of the ways the

set of terms selected for inclusion in a given ontology changes from version to version as a strategy for evaluation of both the ontology as a whole and of the contributions made to this ontology by specific individuals and groups.

Examples of general terms used by scientists that are unproblematically (assuming no errors in the corresponding scientific theories) such as to represent types, include:

(D) Boson, electron, organism, planet, apoptosis, death, orbit.

Examples of general terms which are unproblematically such that they do *not* represent types include:

(E) Thing that has been measured, thing that is either a fly or a music box, organism belonging to the King of Spain, case of pneumonia in man wearing uniform while riding bicycle on small boat with or without fall from stairs.

Note that the terms on either list can be used unproblematically to formulate representations of collections (however the latter term is to be understood³). Terms like those in (D) however, can be used to define collections in accordance with (C) above.

Given entities might be similar along different dimensions. For example they might be similar with respect to *length*, or *feeding pattern*, or *distance from Witwatersrand*. Informally, we can say that, for the collections defined by terms like those in (D) there is a relation of similarity that holds between the members of the collection in virtue of *what they are* (for example *cells*). For the collections defined by terms like those in (E), in contrast, the pertinent similarity relation holds between the relevant members because of *how they are* (for example *how they are related to locations* or *observers*).

Note, however, that in many cases even terms of the latter sort can still be *defined* in terms of types, as for example in:

(F) Thing that has been measured =_{def.} thing that has served as target of some instance of the type *act of measurement*.

This strategy for re-defining terms will turn out to play a central role in our understanding of ontologies in what follows. It can also help to elucidate the relation between *universals* or *types* on the one hand, and *repeatables* on the other. Roughly: wherever we have descriptions of repeatables of the form ‘the X’s’, some way can be found to define the ‘X’ term along the lines of (F) above.

1.5. The Higgs boson

When scientists attempt to detect the Higgs boson (Abazov et al., 2010) they are seeking, first of all, to detect certain particulars – individual things that exist (albeit in some merely probabilistic sense) in space and time. But they are not, of course, seeking to detect just any particulars. Rather, they are seeking particulars that are similar to each other in the sense that they are, again, instances of a corresponding type.

In the case of successful detection, scientists would accordingly need to report their results by employing descriptions of two sorts. On the one hand, they would need to use individual

³The ontology of collections is itself a difficult subject, and we can provide only brief and informal indications here. In a full account we would need to address the question whether phrases like ‘all members in a collection’ mean: ‘all members existing *at a given time*’ or ‘all members existing *at any time*’ (Ceusters & Smith, 2010). We would also need to address the issues of vagueness which arise where similarity relations are marked by gradients (Smith & Brogaard, 2000; Bittner & Smith, 2001). Such issues will not, however, affect our argument here.

referring expressions to identify what they had observed in particular experiments. This would yield sentences such as:

Higgs boson particles have been detected by the CERN Large Hadron Collider in an experiment carried out on June 4, 2014.

On the other hand, they would need to use general nouns and noun phrases to refer to the types whose instantiation had been predicted by the relevant scientific theories, for example, in sentences of the form:

All six types of elementary boson predicted by the Standard Model (photon, W boson, Z boson, gluon, Higgs boson and graviton) have now been experimentally confirmed.

Here again the word ‘type’ is being used to refer to an entity that is *repeatable*. The underlying idea is that where there is repeatability there are entities called types. Because these entities stand to each other in relations of greater and lesser generality, they can sometimes usefully be represented in corresponding hierarchically organized ontologies, as in Fig. 1.

1.6. Reference ontologies

We can now formulate the following

Reference ontology principle: A reference ontology is a regimentation of the terminological content of the settled portions of a given scientific discipline. It includes general terms used by scientists working in that discipline, which are assumed by the developers of the ontology to refer to corresponding types or universals in reality. It also includes assertions of certain relations between instances of the corresponding types.

Note that this principle is for two trivial reasons immune to certain sorts of criticism. First, it is definitional: a reference ontology is *defined* as an ontology that is created in such a way as to serve the representation of what are assumed to be corresponding types in reality. Second, it is a normative principle having the nature of a *conditional recommendation*: if you wish to create what we call a ‘reference ontology’, then you should conceive what you are doing in such and such a way. The principle speaks not at all to those who wish to create information artifacts of other sorts.

Note also that ‘settled’ does not mean: known to be true. Settled science comes closer, rather, to what (Kuhn, 1970) refers to as normal science. It is a result of its goal of remaining in conformity with settled science that the Gene Ontology does not contain the term ‘gene’.

A second methodological principle can now be formulated, in this same normative spirit, as follows:

Principle of consistency with established science: The assertions of which a reference ontology consists at any given stage should be consistent with the best available settled science that is current at that stage.

The two mentioned principles might in theory be consistent with an approach according to which ontology developers working in support of different scientific disciplines would develop representations of the types in the corresponding domains according to their own specific ideas of how such a task might best be realized. Some might for example decide to create a mere list of types organized alphabetically. Others might create a representation of types organized hierarchically according to the mereological relations between their instances. An uncoordinated approach along these lines would not, however, address the goal of cross-disciplinary data integration. Where neighboring scientific disciplines are formulating results concerning the entities in areas where their domains overlap, we need to

ensure that two ontologies agree in the ways these types are represented. Thus where one discipline deals with subtypes of types falling within the purview of another discipline, then the former will need to classify these subtypes by using terms taken over from the latter.

To address such issues, the representations of types created to support the integration of data generated by a given family of scientific disciplines – for example, in biomedicine – need to be developed in a highly constrained way in conformity with certain common principles that are accepted in advance by the developers involved. The latter will need to agree not only in use of terms, but also in definitions, and this will bring the need for common principles concerning how terms are to be defined. Ontologists will need to agree also in the logics used for reasoning with these definitions, on practices for use of identifiers, for versioning and obsoleting, and for use of ontologies in annotations – and all of this will require a further layer of principles relating to governance and to testing and selection.

1.7. Ontology path dependence

What the various principles should be that guide ontology development is of course the 666 dollar question of ontology coordination. As we shall argue below, to have any hope of success in an area as broad as the entirety of the life sciences, the principles must be understood as part of an evolving, empirically guided process beginning with initial formulations that address as closely as possible readily identifiable needs and practices of biologists, moving on from there in stages to progressively more rigorous formulations allowing incrementally more ambitious approaches to the integration of data.

Our experience tells us that the needed set of principles will involve some which take the form of substantive or technical guidelines for building ontologies (for example, *distinguish continuants from occurrents*; employ a backbone is_a hierarchy using single inheritance). Some principles, however, will be a matter of social coordination, the most important of these being:

Ontology path dependence principle: The decisions made by the creators of an ontology – including those decisions which pertain to the ontology's upper-level architecture – should as far as possible be made on the basis of the degree to which they advance the consistency of that ontology with the reference ontologies already existing in relevant domains.

One of Merrill's central criticisms relates to our acceptance of what he calls the 'Referential Assumption' (Merrill, 2010, p. 85), which (in simple terms) we can express as the proposition that ontologies should consist of general terms as their representational units. Merrill criticizes our work because (he thinks) we hold this belief for complicated philosophical reasons, which he rightly sees as being irrelevant to the practical purposes of science. His criticism is however undermined because he fails to take account of the degree to which we take path dependence seriously (because not to do so would doom our project to failure). Thus he does not comprehend that, for us, the thesis that ontology developers should focus on general terms when constructing ontologies is to be recommended for the simple reason that all successful ontologies in support of science created thus far consist overwhelmingly of representational units of this sort.

Tacit acceptance of the ontology path dependence principle among our biologist colleagues has brought it about that certain ontologies in the area of biology – in particular the Gene Ontology (GO) – have come to enjoy a privileged position. The GO is a controlled vocabulary developed to serve the consistent formulation of information pertaining to the *attributes of gene products* in organisms of different types (Gene Ontology Consortium, 2000). Since its creation in 1999 it has enjoyed a phenomenal success, and its role as *de facto* standard ontology in important areas of biology makes it in some ways comparable to

the US interstate highway system. This in turn justifies the expenditure of extraordinary effort to ensure that it continues to be developed in ways that maintain its consistency with the best available science.

This privilege reflects in part a simple *homesteader effect*; since ontology is so new, there are many fields thus far not ontologically tilled. The first in the field in any given area acquires certain presumptive rights. One such right consists in the fact that ontologies developed thereafter in neighboring domains have a responsibility to ensure that they are constructed in ways that make them consistent – from the point of view of both logico-ontological architecture and scientific content – with the already privileged ontologies which came earlier. In addition, it implies that certain design choices made in the construction of these established ontologies should, again presumptively, be adhered to also by the successor ontologies which are created in their wake. At the same time, of course, the homesteader privilege brings considerable responsibilities, and the presumptive rights associated therewith can in principle be over-turned in case of demonstrably poor husbandry (Smith & Ceusters, 2006; Smith, 2006b; Smith, 2010).

1.8. Asserted monohierarchies

Inspired in part by (Rector, 2003), we advocate the following:

Principle of asserted single inheritance. Each reference ontology module should be built as an asserted monohierarchy (a hierarchy in which each term has at most one parent).

This means that the ontology will have a single root node, and that all non-root terms will have exactly one *is_a* parent and thus be connected by exactly one chain of *is_a* relations to the root. To say that the *is_a* relations are *asserted* means that they are included in the ontology manually by the ontology's developers and form the basis for the associated definitions.

Terms in the resultant asserted hierarchies can be used in various combinations, using relations taken over from the Relation Ontology (RO) (Smith, 2005) to form new terms, following a methodology first applied in relation to the GO and its sister ontologies in Wroe et al. (2003) (compare also Hill et al., 2002; Mungall, 2004). The goal is both to reduce the degree of arbitrariness typically involved in term composition in ontologies, and to ensure that ontologies are developed in tandem in such a way as to constitute a progressively more well-integrated modular network.⁴ A term such as *blood glucose measurement*, for example, is formed from FMA:*portion of blood*, ChEBI:*glucose* and OBI:*act of measurement*. When a classifier is applied to the result of adding such a term, with its definition, to the already existing set of asserted monohierarchies, then certain further *is_a* relations will be able to be *inferred*. This will in some cases yield a polyhierarchy, or in other words a hierarchy in which some terms will have more than one *is_a* parent (hence 'multiple inheritance' – meaning that the entities represented by a term with multiple parents will *inherit* a corresponding set of attributes from each of its parents).

Rector (2003) has developed a methodology for 'normalizing' ontologies by decomposing existing polyhierarchies into homogeneous disjoint monohierarchies. For him, the monohierarchies are then recombined using logical definitions from which an enriched polyhierarchy can be inferred mechanically using a theorem prover or reasoner.

⁴http://www.berkeleybop.org/ontologies/#logical_definitions, last accessed June 30, 2010.

We welcome *post hoc* applications of Rector's normalization process where ontologies – as in the case of the GO – exist already in a non-normalized state. We go further, however, in advocating the formulation of new reference ontologies *ab initio* as asserted monohierarchies, and multiple such ontologies have been created already within the context of the OBO Foundry initiative.

The goal from our point of view is that the resultant normalized ontology modules should as far as possible reflect the existing disciplinary division of labor in the relevant domain of science. Relevant *inferred* polyhierarchies can then be created according to need, for example, when providing support for information retrieval or for the representation of multi-disciplinary scientific content or of the results of a particular set of experiments. The approach helps to ensure that each ontology module is practically surveyable, thereby supporting the purposes of effective maintenance and use; but it also supports more effective computation, since it is easier to write software for normalized ontologies (Rector, 2003). It is also easier to formulate, to explain and to understand ontology definitions formulated in terms of types and attributes than in terms of multiple subsumption relations, so that the approach serves also to ensure consistency between formal and natural language definitions, and in this way it allows human maintenance of the definitions in the ontology to be carried out in tandem with consistency checking via software.

As Rector shows, ontology-based integration is easier to manage and scale on the basis of normalized ontology modules. It is easier to master the problems associated with combinatorial explosions when normalized ontology modules and a restricted set of relations are used to serve as the basis for allowable sorts of combinations. It is also easier to maintain ontologies, for example, when a change must be made due to some scientific advance. This is because the change in question can be made in just one place in the normalized ontology, allowing consequent changes in the associated polyhierarchies to be propagated automatically.

As we ourselves have argued at length (for example, in Smith, 2004), ontologies which allow multiple inheritance are prone to characteristic kinds of errors, not least because different axes of classification become hard to keep separate in developers' minds. And as Rector points out, 'Empirically, whenever examining a multi-axial "ontology" and then normalising it, we find errors'.⁵ We have also seen in our own experience of working in ontology teams with, for example, plant, or cell, or infectious disease biologists, that the restriction to single inheritance, while often initially painful because it is seen as placing restrictions on what can be said, very often yields a solution that is seen by the developers as more illuminating of the underlying science, and thus as more stable, than the multiple inheritance-based approaches that had previously been adopted.

Normalized ontology modules help in preventing errors also because their plug-and-play character helps to encourage ontology reuse. Because those who are called upon to construct new ontologies are more easily able to draw upon ontology content that has already been thoroughly tested, they therefore do not need to construct ontology components anew and thus they can avoid creating new errors and inconsistencies, and thus new avenues for silo formation.

1.9. Basic Formal Ontology (BFO)

Restricting the asserted portions of the *is_a* hierarchies of reference ontologies to single inheritance thus brings considerable benefits, and – because there is an easy way of then

⁵http://obofoundry.org/wiki/index.php/Asserted_Single_Inheritance, last accessed August 7, 2010.

generating the associated multiple inheritance-based artifacts people might need – these benefits come at very little cost. Some in the GO community are accordingly proposing experimentally to figure out what the normalized versions of the three Gene Ontologies would have to be in order to ensure that the existing versions of the ontologies could be derived automatically therefrom by using reasoners. Even a partial success in this regard would add much to GO's utility to reasoning systems.

In the case of the GO, it is clear what the relevant root nodes should be in such a normalized reconstruction – they would be *cellular component*, *molecular function*, and *biological process*, respectively, corresponding to the three existing divisions of the GO. In the general case, therefore, it is not so clear how such root nodes for normalized ontology modules should be selected and how they should be positioned in relation to the root nodes of neighboring ontologies. If ontologies are to be developed in coordinated fashion, therefore, then substantive principles need to be available also to support the making of decisions such as this, and to this end we need a strategy concerning which most general types or universals should be taken as the starting point for the process of populating an ontology downward from the root. To this end we have proposed the set of categories together forming the Basic Formal Ontology (BFO) (Grenon and Smith, 2004), specifically the three top-level categories of *independent continuant*, *dependent continuant* and *occurrent*.

Some set of upper-level categories is needed if ontology coordination in the service of data integration is to be possible at all, and in Section 6 we shall argue the merits in this regard of BFO resulting from the fact that it was created for precisely this purpose. Already some 75 ontology projects in different domains of the life sciences are being developed in its terms. The authors of the Foundational Model of Anatomy have for some years been working to ensure conformance with BFO (Smith & Rosse, 2004; Rosse & Mejino, 2007). BFO has also been subjected to thorough tests of its serviceability as an upper level ontology for scientific purposes by the members of the OBI (Ontology for Biomedical Investigations) Consortium, and by users and critics such as Thomas Bittner, Maureen Donnelly and Randall Dipert (Buffalo), Mathias Brochhausen (IFOMIS), Lawrence Hunter and Mike Bada (Denver), Chris Mungall (Berkeley), Fabian Neuhaus (NIST), Bjoern Peters (San Diego), Alan Ruttenberg (Buffalo), Holger Stenzhorn (IFOMIS, Saarland University) and Kerry Trentelman (Buffalo), as well as by some of the 120 members of the BFO Discussion Group.⁶ These tests have led to a number of changes in the ontology over time. They have also, as we are the first to admit, identified a number of shortcomings in BFO, some of which (we hope) will be addressed in the forthcoming release of BFO 2.0.

One important feature of BFO is that its tripartite top-level structure echoes the tripartite design of the Gene Ontology. Collaboration between the BFO and GO Communities was inaugurated at a meeting organized in Leipzig in 2004 on the topic of *The Formal Architecture of the Gene Ontology*,⁷ where Smith, in a presentation entitled “STOP!”⁸ presented arguments in favor of the need for certain changes in the GO. These arguments received a favorable response from the GO Consortium because they were seen as bringing immediate practical benefits, including:

- a. providing a clearer understanding of the relation between terms in the GO and the entities studied in biological experiments (Hill et al., 2008),

⁶<http://groups.google.com/group/bfo-discuss/members>, last accessed June 30, 2010.

⁷http://www.ifomis.org/Events/GenOntology_2004/, last accessed June 30, 2010.

⁸For: Smart Terminologies through Ontological Principles.

- b. providing a readily applicable technique for formulating definitions of the *is_a* and *part_of* relations and thereby removing certain inconsistencies in GO's earlier treatments (Smith, 2004),
- c. identifying errors in terms and definitions of GO, leading for example to the obsolescence of terms such as GO:0005941 *unlocalized protein complex*, which reflected a confusion of ontology with epistemology.

One result of our work with Ashburner, Lewis, Lomax, Mungall and other GO principals, and also with the leaders of the FMA and GALEN groups, was the creation of the Relation Ontology (RO) (Smith, 2005), which is designed to restrict the repertoire of relations available for use by biomedical ontology developers to a small set, all the members of which are logically defined in such a way as to promote interoperability of the ontologies which use them.

Following shortly after the publication of the RO paper came the establishment, in 2006, of the OBO Foundry (Smith et al., 2007), which adds a layer of governance and of peer review to the process of multi-ontology development, and which uses the GO/BFO tripartite division of categories as basis for partitioning the totality of biomedical entities into non-overlapping ontology domains (see Fig. 2).

1.10. How the ontological realist methodology works to support ontology authoring

Our methodology for ontology development requires that discipline-specific reference ontologies be created manually by experts in the corresponding disciplines, persons who already know what it is in reality to which the terms in their discipline refer. The first round in the iterative process of building a discipline-specific ontology will require the creation by such persons of a draft list of the general terms that can be used within the discipline in positive assertions to refer – on initial inspection – to types or universals.

For any given settled science the set of candidate terms in this respect is broadly understood and accepted by the scientists involved. The problem is that this set is typically too large for the purposes of coordinated ontology development. Some terms will thus need to be removed, for example because of redundancy or ambiguity, or because they refer not to a corresponding universal or type, but rather to what we might refer to as an attributive collection of particulars, as for example, 'human who has been tested for HIV' or 'human with bra cup size C'.⁹ Further terms, such as 'known allergy', 'other diabetes', 'pneumonia diagnosed by inspection of sputum sample', will need to be excluded because they involve a more or less hidden reference not to the way things (repeatably) are on the side of reality but rather to some particular feature of our present state of knowledge (Bodenreider et al., 2004).

To ensure conformity to the principle of asserted single inheritance, it will sometimes be necessary to transplant some terms from the initial list into separate lists, for example, by following the recommendations generated by application of the OntoClean methodology (Guarino & Welty, 2002). The transposed terms will then be defined using terms which remain, together with terms from other reference ontologies according to need. In this way, for example, a term such as 'mechanosensory organ' might be removed from a structurally based anatomy ontology and defined in terms of the anatomical term 'organ' and a term such as 'mechanosensory function' created in an external function ontology. The

⁹In Smith and Ceusters (2006) we called such collections 'defined classes'. We no longer favor this terminology since the fact that a given term is or is not defined in a given ontology need carry no significance as to the status or nature of the entity represented.

transplanted terms rest, in effect, on classificatory principles skew to those adopted in building the ontology with which one begins.¹⁰

When the asserted monohierarchies have been identified, the terms in each hierarchy can be defined according to the

Principle of Aristotelian definitions (Rosse & Mejino, 2003): Given a term ‘A’ in an asserted monohierarchy, with parent term ‘B’, the definition of ‘A’ should take the form

$$A =_{\text{def. a}} B \text{ which } C's$$

where ‘C’ expresses some condition on those instances of B which fall within the A’s.

One consequence of this principle is that there are no disjunctive or conjunctive or negative universals – an issue to which we return in our treatment of the term ‘non-smoker’ below.

1.11. The principle of instantiation

The inclusion of a representation of a universal in the GO requires that at least one real-world instance of this universal has been shown experimentally to have existed. Consider, for example, the universal *retinol dehydrogenase activity*, defined as the potential to realize the reaction: $\text{retinol} + \text{NAD}^+ = \text{retinal} + \text{NADH} + \text{H}^+$. Before this term could be included in the GO’s molecular function ontology, it was necessary that experimental evidence be provided (Zhang et al., 2001) to the effect that there exist molecules that have instances of this universal as their functions.

GO’s practice here is taken as model for a further principle by means of which ontology authors can judge whether given terms should be included in a given ontology, namely the

Principle of instantiation: A term should be included in a reference ontology only if there is experimental evidence that instances to which that term refers exist in reality.

(‘Exists’ here should be understood in a tenseless sense in order to accommodate, for example, universals pertaining to extinct species as well as universals such as *swarm* or *hurricane* which are instantiated only intermittently.)

Insisting upon the principle of instantiation, and thus on experimental evidence, provides us also with a means by which we can judge whether two ontologies are orthogonal in the sense that they do not overlap in their respective domains. Is an ontology containing the term ‘phosphogluconate pathway’ orthogonal to an ontology containing the term ‘pentose phosphate cycle’? To find out, we need to identify what types (if any) these terms refer to in reality, and for this we will need to work with biologists who are carrying out salient experiments and who can thus explain to us what processes of intervention and observation are involved in gaining information about the corresponding instances.

References to different sorts of universals are in this way used to form chains of validation, whereby tests for the instantiation of universals further down the chain (for instance, *molecule of retinal*) provide evidence for the existence of universals further up (which means: universals, or putative universals, closer to the frontiers of current knowledge – for instance, *retinol dehydrogenase activity*). Often very simple universals are involved in such validations, as for example when instances of the color universals *purple*, *pink* and *red* are

¹⁰Sometimes there will be several ways of achieving the end of single inheritance. (Compare the situation in topology where any one of a number of basic terms such as ‘boundary’, ‘closure’, ‘interior’, ‘open’, ‘closed’ can be selected as primitive in such a way that each of the other terms on the list can be defined therefrom.)

observed inhering in instances of *sputum* in applications of the Gram staining protocol, thereby allowing inferences to the effect that instances of given types of *bacteria* exist in tested samples.

An analogous scenario applies at the level of instances. A clinician, for example, has observed rales and rhonchi upon examining a patient, and hypothesizes that she is suffering from pneumonia. To verify this hypothesis the clinician does not look for an instance of the universal *pneumonia* that he would somehow be able to observe directly. Rather, he looks for instances of certain more easily confirmable universals, for example, by using sputum tests to determine the presence of instances of *blood*, of (*Gram-positive* or *Gram-negative*) *bacteria*, of *antibodies to S. pneumoniae* and so on.

1.12. A system of reference ontologies

When once a somewhat stable set of normalized *is_a* hierarchies has been created, the terms in the resultant graph-theoretic structures need to be linked further by relations of other sorts. Such links will need to be established both to other terms in the ontologies being developed and to relevant terms in neighboring ontologies.

Our modular strategy rests hereby on a division of labor between ontologists in different disciplinary communities working in tandem on the basis of BFO as common formal ontology, the latter being itself subject to revision in light of its ability to serve the representations of the corresponding portions of science. The need for an approach involving a common upper level ontology is, we believe, a simple practical consequence of collaborative ontology development in the service of empirical science. The presence of a common upper ontology means, for example, that those working on cells or proteins are easily able to draw on each others' resources in building their respective ontologies, revising these in tandem in reflection of changes brought by advances in empirical science (Masci et al., 2009). All of those involved are thereby engaged in creating not merely the ontologies themselves but also, as an inevitable side-effect, an evolving set of mutually binding constraints on each others' work that serves to ensure that these ontologies are developed in such a way that their interoperability is preserved over time. These constraints (principles, criteria) must be widely acceptable to different groups of scientists providing data for integration. At the same time, they must be able to bring about a process of evidence-driven improvement in the ontologies constructed in their terms. The result is a system of reference ontologies whereby:

1. for any given domain of reality, exactly one *reference ontology* is constructed that is (a) in conformity with the settled science in that domain and (b) capable of being recommended for general use,
2. these orthogonal reference ontologies will be semantically interoperable with one another,
3. they will reduce the need for (typically fragile and costly) mappings between ontologies covering the same or overlapping domains, and
4. they will be able to be used as a reliable starting point for the development of application ontologies needed for specific purposes.

The ontological realist methodology has been embraced by a growing number of researchers, some of them leaders in their respective fields, and we and others have devoted considerable efforts to refining the methodology and disseminating its principles among a variety of different biologist, clinician and informatician communities. Important users include, in addition to those listed in Section 1.9, also Cornelius Rosse, Melissa Haendel, Onard Mejino in the FMA and CARO anatomy ontology projects (Rosse & Mejino, 2007;

Haendel et al., 2008); Melanie Courtot, Philippe Rocca-Serra, Susanna-Assunta Sansone, Chris Stoeckert, and the late Bill Bug in the Ontology for Biomedical Investigations Consortium¹¹; Lindsay Cowell, Alexander Diehl, Albert Goldfain, Yongqun He, Anna Masci, Kitsos Louis, Richard Scheuermann and their colleagues in the Infectious Disease (Topalis et al., 2010) and Cell Ontology Consortia (Diehl et al., 2010); Maryann Martone and her colleagues developing the Neuroinformatics Information Framework (NIF) Standard¹²; Cecilia Arighi, Judith Blake, Darren Natale, Cathy Wu, in the Protein Ontology Consortium (Arighi et al., 2009); Colin Batchelor, Karen Eilbeck, Janna Hastings and Neocles Leontis in the Sequence Ontology (Mungall et al., 2010), CHEBI ontology of small molecules (de Matos et al., 2010) and RNA Ontology Consortia (Batchelor et al., 2009); Ramona Walls, Laurel Cooper, Dennis Stevenson, and Pankaj Jaiswal of the Plant Ontology Consortium¹³; and Sivaram Arabandi, Albert Goldfain and William Hogan of the OGMS Ontology for General Medical Science initiative¹⁴ – as well as friendly disposed observers such as Stefan Schulz (Schulz et al., 2009), Kent Spackman (Ceusters, 2007), Olivier Bodenreider (2008) and Georges De Moor, founder of CEN/TC251 (Ceusters et al., 2009). The intellectual firepower of these authors and of their collaborators, together with that of Michael Ashburner, Suzanne Lewis, David Hill, Jane Lomax and their Gene Ontology colleagues, has contributed immensely to the content of this methodology and to its progressive refinement.

Each reference ontology, if our strategy is successful, will, like the GO, serve as an attractor for multiple expanding groups of users whose members will have strong incentives not only to invest resources directed toward ensuring that it is developed and used in ways that keep pace with scientific advance, but also to recommend it to other users – since this will increase the value of their own investment. In this way, we believe, we have a strategy which can avoid recreating through ontology proliferation the very silo effects to which ontologies themselves were originally conceived as the antidote (Smith, 2008). We know of no other approach to ontology development of which an analogous claim can be made.

1.13. How the ontological realist methodology works to support ontology maintenance

Scientists in many areas of biology, including clinical research, have come increasingly to rely on a process whereby professional biocurators manually create *annotations* to experimental data using terms from the GO. This annotation process unfolds in a series of steps, which can be summarized as follows:

1. the curator identifies specific experiments documented in the scientific literature in which instances of (for example) specific types of protein interaction have been detected in observations;
2. the curator applies expert knowledge to the documentation of the results of these experiments, a process which involves determining which types of gene products are being studied in the experiment, and which types of molecular functions, biological processes and cellular components are identified as being correlated therewith;
3. the curator creates an annotation, which captures the relationships between identified gene product types and the corresponding Gene Ontology types, and which is then added to an annotation database;

¹¹<http://obi-ontology.org/page/Consortium>, last accessed June 30, 2010.

¹²<http://ontology.neuinfo.org/NIF/nif.owl>, last accessed June 30, 2010.

¹³<http://www.plantontology.org/>, last accessed June 30, 2010.

¹⁴<http://code.google.com/p/ogms/>, last accessed June 30, 2010.

4. where representations of specific types needed for annotations are missing from or misclassified in the GO, the curator submits a corresponding request for inclusion or correction to the ontology's editors using a dedicated tracker.

Through the implementation of step (4), a virtuous cycle is brought into play in conformity with what we shall call the:

User feedback principle: A reference ontology should evolve on the basis of feedback derived from those who are using the ontology for purposes of annotation.

This means that the process of curation of experimental results by biologists contributes to the on-going improvement of the ontology. This in turn contributes to improvements in the annotations created in subsequent cycles.

The methodology is described in detail in Hill et al. (2008), which makes clear the essential interplay between the two kinds of descriptions referred to already above of (i) the individual entities observed in the lab and captured in reports of experiments, and (ii) the types these entities instantiate, which are represented through the use of general terms in the assertions of the corresponding scientific theories.

The idea underlying our methodology for the development of such reference ontologies can now be summarized as follows. Scientists formulate assertions describing their experimental results and publish them in scientific papers and textbooks. These assertions contain expressions of various sorts, some of which are candidate referring expressions. Some of the latter will be general terms specific to the discipline in question, expressions used by scientists to formulate assertions with positive intentional force, such as 'Bosons are particles which obey Bose–Einstein statistics' or 'The N-terminus of retinol dehydrogenase type 1 signals cytosolic orientation in the microsomal membrane'. When initiating the development of a reference ontology for a given scientific domain, we adopt, for each term used by the given science, a defeasible assumption to the effect that it refers to some corresponding type or universal. This assumption can be overturned in a number of ways. Most interestingly, it can be overturned by scientific discovery, as for example, in the case of 'phlogiston'. By default, however, the assumption holds simply because as soon as it becomes known to the scientists involved that a given general term does not refer to any corresponding type or universal, then this term will be dropped from the repertoire of those terms that can be used in the normal assertive contexts of the relevant science.

1.14. How can we know that a given general term denotes a universal?

Merrill's own approach to matters terminological awards a central place to *meanings* (Merrill, 2009). Trying to figure out what words or phrases mean in source vocabularies is indeed the primary duty of, for example, the curators of the UMLS Metathesaurus. It is not, however, what ontological realism is all about. The ontological realist methodology for building ontologies starts, rather, with the *terms* used by scientists and with the *particulars in reality* that these terms are used to describe. The issue of ambiguous terms that is so important for the approaches studied in (Merrill, 2009) is nipped in the bud at the very start of the development process for each ontology by excluding or appropriately relabeling terms that are used by scientists in ambiguous ways.

Our writings on the ontological realist methodology set forth principles which describe how to create ontologies in such a way that any reader who is familiar with these principles and with the relevant science (including the relevant types of scientific experiment) can know exactly what is intended to be denoted by the terms the ontologies contain. At no point do we make any appeal to meanings. At no point do we refer to relations of synonymy. And at no point other than when addressing views proposed by others do we refer to concepts.

The principles are codified in our Referent Tracking (RT) framework (Ceusters & Smith, 2006b), only one element of which is discussed by Merrill – namely, what we call the ‘PtoU-tuple’ template (for ‘particular-to-universal’) (Merrill, 2010, p. 96). Examining his remarks in this connection will make clear why our proposals cause him such consternation and why, on the basis of a proper understanding, this consternation could have been avoided.

The PtoU-tuple template pertains to the RT-recommended syntactic regimentation of a statement, authored by a particular a , to the effect that some universal u , referred to in some ontology o , is instantiated by some particular p :

$$\langle IUI_{a;t_a}; \mathbf{inst}; o; IUI_{p;u;t_r} \rangle.$$

Here ‘ IUI ’ stands for ‘instance unique identifier’. When this template is used to create an actual tuple that is intended to describe some portion of reality in an RT-conformant fashion, then ‘ u ’ is replaced by the designation, taken from some pre-existing ontology, of some universal with which the particular denoted by ‘ IUI_p ’ enjoys the instantiation relationship (**inst**). As we explain at length in Ceusters and Smith (2006b), the PtoU-template is introduced precisely to express the instantiation of some universal by some particular. If John Doe, a follower of ontological realism, formulates a statement by means of this template, for example along the lines of:

(G) \langle John Doe; 06/11/2010:6.45PM; **inst**; BFO; Barry Smith; independent continuant; since 1951 \rangle

then this implies that John Doe believes the following:

1. that Barry Smith and BFO are particulars,
2. that BFO is an ontology,
3. that ‘independent continuant’ denotes a universal,
4. that the instantiation relationship between Barry Smith and the universal called ‘independent continuant’ has obtained since 1951.

John Doe might be wrong in one or more of these beliefs, and in that case his statement (G) is false.

Merrill now expostulates as follows: ‘ u in such an entry is said to be the name of a universal. Now why should we suppose that it is?’ (Merrill, 2010, p. 96) This is, from our perspective, a bit like hearing someone responding to the assertion: ‘the right to a speedy and public trial is one of the rights enumerated in the Constitution of the United States’ by saying: ‘Now why should we suppose that it is?’

Ontologies from our hand contain representational units that are assumed to denote universals or types in reality. (Recall the *reference ontology principle* in Section 1.6.) That is how, in the context of the referent tracking literature Merrill is here criticizing, an ontology is *defined*. Data repositories that follow the referent tracking paradigm similarly contain exclusively individual identifiers that are intended to refer to particulars. The underlying idea can be codified in the form of a principle now observed successfully for some years by ontologies such as the Gene Ontology:

Principle of obsolescence: Should we ever find that a term in an ontology or data repository fails in designation, then the relevant entry will immediately be obsoleted. This applies to expressions referring both to what is general and to what is particular. (Ceusters, 2007; Ceusters & Smith, 2006a.)

2. The background of ontological realism

2.1. Why ‘universal’?

We have used the word ‘type’ in the above side by side with the word ‘universal’. The use of ‘type’ reflects an effort on our part to be responsive to the needs of specific communities of readers. But it has at the same time caused confusion because ‘type’ is used in multiple different ways in the multiple disciplines relevant to ontology. The term ‘universal’, in contrast, has an established narrowly defined use that serves our ontological realist purposes very well, and it is for this reason that we employ this term as part of our technical vocabulary.

One downside arising from the choice of a term of such ancient provenance is that its usage sets certain sorts of philosophically trained individuals into something approaching panic. (This is true, with especial potency, in the case of Merrill (2010, p. 93) – whose *shock* at the fact that, still today, someone might use this word in a serious way – is tempered only by the fact that he himself employs with a similar purpose the term ‘category’,¹⁵ a term likewise deriving from Aristotle.)

The countervailing benefit we derive from using ‘universal’, however, is that the term conforms to the

Minimal terminological baggage principle: When working in a multidisciplinary field such as ontology, avoid the use of technical terms that have multiple conflicting technical uses in the constituent disciplines involved.

It is with the aim of conforming to this principle that we try to avoid in our presentations of the ontological realist methodology also terms such as ‘class’, ‘property’, ‘model’, ‘semantics’, ‘thing’ and of course ‘concept’ (Smith, 2004). In particular we prefer to use ‘collection’ rather than ‘class’ in order to avoid the drawing of conclusions by developers of ontologies from doctrines pertaining to the use of the word ‘class’ in Description Logic contexts.

2.2. Universals in language and cognition

Our own ideas on universals derive from our study of the work of Edmund Husserl, whose *Logical Investigations* contains the first use of the term ‘formal ontology’ (Husserl, 1913/21, II, p. 219, 1970, pp. 428f). Husserl describes certain universal laws governing how parts are related within structured wholes, laws, for example, of the form: if an instance of the universal *A* exists within a given whole, then so also will an instance of a second universal *B* (Smith, 1987; compare Smith et al., 2005). Simple examples are found in perceptual psychology: every sensation of color involves some sensation of visual extent. But it was in the field of linguistics that Husserl’s ideas were particularly influential, where they led to the creation of what is now called ‘categorial grammar’ (Buszkowski et al., 1988), and where they influenced also the work of structural linguists such as Jakobson (Holenstein, 1976), of the early speech act theorists (Smith, 1990), as well as Chomsky’s idea of a ‘universal grammar’ (Kuroda, 1997).

Parallel developments in linguistics led also to the work on *universals of human language* on the part of Joseph H. Greenberg and his followers (Greenberg, 1963). In Greenberg’s terms, all languages have *nouns* and *verbs* and all spoken languages have *consonants* and *vowels*. What, he asks, are the other universals common in this way to all human languages? The attempt to answer this and a series of analogous questions initiated what is still one of

¹⁵<http://www.ncsu.edu/chass/philo/LACSI.Abstract.pdf>; <http://biometrics.com/wp-content/uploads/2009/06/safetyworks.pdf>.

the most powerful research programs in the cognitive sciences. The project has been influential also in disciplines such as anthropology, for example, in Brown (1991), which identifies some hundreds of cognitive and behavioral universals common to all human societies (compare also Pinker, 2002). And because the evolution of languages is influenced by the same population splits that influence human genetic changes, work on language universals has provided valuable materials also in assisting population geneticists trying to reconstruct the path of early human migrations by means of genetic patterning in different peoples (Cavalli-Sforza, 1997).

Interestingly, Greenberg's work on universals and on the typology of language grew out of his deep study of Aristotle, and he followed Aristotle's empirical methodology for identifying universals through inspection of many examples.¹⁶ The reader may thus be wondering if the world has reason to be grateful for the fact that the successes of Greenberg and his followers in throwing light on human cognition and behavior were not thwarted by complaints, from some Merrill counterpart of an earlier era, to the effect that they were associating themselves with a metaphysical tradition with a 'long and sordid history' (Merrill, 2009, note 8).

2.3. Universals, scientific realism and received first-order logic

In an independent development in the late 1970s the term 'universal' began to be used by philosophers as part of a general rediscovery of the importance of traditional metaphysical thinking, and especially of one or other version of metaphysical realism, for an understanding of scientific laws. This rediscovery occurred after a period of dominance of nominalism especially among philosophers active in the United States who were taking advantage of the possibilities created by the new tool of first-order predicate logic (FOL) for the formulation of philosophical arguments.

Simply put, the formulae of FOL consist of four kinds of expressions: logical constants, such as 'and' and 'not'; quantifiers such as 'all' and 'some'; constant and variable terms such as '*a*', '*b*', '*x*', '*y*'; and predicates such as 'F' and 'R'. Formulae such as 'F(*a*)' or 'R(*a*, *b*)' are then used to regiment natural language assertions such as, respectively, 'Socrates is a man' and 'Socrates is married to Xanthippe', where '*a*' stands in for 'Socrates', '*b*' for 'Xanthippe', 'F' for 'is a man' and 'R' for 'is married to'.

Fatefully, Quine and some of his contemporaries succeeded in establishing a widespread presumption according to which the use of FOL as a tool of philosophy must go hand in hand with the acceptance of a rather narrow (and nominalist) view as concerns the range of entities to which constituent *terms* of FOL are allowed to refer. Specifically, the view came to be adopted according to which all terms in FOL must refer exclusively to individual objects (particles, molecules, cells, organisms, planets and so forth). The result – which we shall henceforth call *received FOL* – reflects, as we shall see, a genuine restriction on the available expressive resources of first-order logic. Yet its influence has been so great that even those thinkers who embraced the new metaphysical turn in philosophy in the 1970s, including David Armstrong, continued to fall victim to it (Smith, 2005).

One reason why Merrill has such a problem with universals, we believe, is because they fall outside the scope of what can be referred to within the framework of received FOL, a

¹⁶As Greenberg's disciple Givón puts it:

You want a theory of universals? Invest some time in looking at a decent sample of diversity. For your theory must account for both. This is how Aristotle assembled his functionalist biology in *De Partibus Animalium* and *De Generatione* – by studying the vast diversity of extant species. This is also the way he came by his political theory in *The Politics* and *The Constitution of Athens* – by poring over scores of diverse extant constitutions. (Givón, 2002.)

framework which, like many 20th century analytic philosophers, Merrill views as the benchmark of acceptable formalization (2010, note 17). Because terms in received FOL range exclusively over individual objects such as molecules or cells or people, such terms cannot be used to refer to universals, or to anything general or repeatable. And the predicates in FOL cannot be used to refer to such entities either – because they cannot be used to *refer* to anything at all.

The metaphysical turn of the 1970s consolidated itself in a new subdiscipline called ‘analytical metaphysics’, which has since become an established part of the philosophical mainstream. The doctrine of nominalism is indeed still alive in some circles of analytical metaphysics today. In a survey of (primarily Anglosaxophone, analytical) philosophy faculty carried out in November 2009, however (Bourget & Chalmers, 2009), only 15.1% of the 931 faculty surveyed described themselves as accepting nominalism.¹⁷

Accordingly, when Merrill asserts that ‘universals, and Aristotelian realism have come under a series of sustained attacks for at least centuries, if not millennia’ (Merrill, 2010), then the reader should be aware that these nominalist attacks have been launched so often precisely because of the remarkable tenacity of the metaphysical realist position.

2.4. Summary of Merrill’s argument

I know that you believe that you understood what you think I said, but I am not sure you realize that what you heard is not what I meant.

Robert McCloskey, State Department spokesman.¹⁸

Merrill (2010) points to some of the reasons why our methodological views have been found attractive by researchers in different life science domains because of the practical advantages they bring to the developers of ontologies. But at the same time, he assails this methodology on a number of grounds, some of which rest on misinterpretations of our views, some of which are, we confess, a consequence of the fact that our position is not easy to grasp from the multiple expositions that we have created over the years for different audiences of users. Ontology, when practiced seriously, is of its nature a multidisciplinary affair, and we believe that our approach has gained traction in no small part because we have taken its different disciplinary dimensions seriously. (Thus we have not viewed ontology as an activity performed by and in the service of, for example, lexicographers, who tend to see ontologies as focused primarily on meanings; we have also not viewed ontologies as the idealized algebraic structures that are of special interest to some computer scientists.) The interactions with these multiple disciplinary groups of users have led also, over time, to important changes in our approach, including changes in our terminology. We are thus grateful to Merrill for having provided us with the opportunity to address some of the misunderstandings resulting herefrom.

Merrill’s own major misunderstandings of our view can be summarized as follows:

1. that ontological realism as we understand it is a metaphysical realism of the sort defended by the philosopher David Armstrong (Merrill, 2010, p. 104);
2. that we hold that studying and embracing metaphysical realism is a requirement for doing science (Merrill, 2010, p. 103);

¹⁷http://philpapers.org/surveys/results.pl?affil=Target+faculty&areas0=0&areas_max=1&grain=fine. 45% of respondents listed Aristotle as the non-living philosopher with whom they most identified (this albeit in some cases for reasons not exclusively metaphysical).

¹⁸(attributed) http://www.quotationspage.com/quotes/Robert_McCloskey/, last accessed June 30, 2010.

3. that we accept what Merrill calls the ‘Referential Assumption’ according to which the so-called general terms of our language (such as ‘man’) participate in a direct reference relation in precisely the same manner as do the singular terms of our language (such as ‘Socrates’) (Merrill, 2010, p. 85).

Under (1), Merrill fails to do justice to the narrowly practical significance of and justification for our proposals. His misinterpretation of our views from this perspective can be summarized as follows: that he interprets a *methodology* recommended for use by ontologists working in scientific domains as a *theory* about the nature of science as a whole. Sometimes such misinterpretation involves creative misquotation on Merrill’s part, as when, for instance, our statement in Smith (2004) to the effect that:

good modeling in support of the natural sciences *can* ... be advanced by the cultivation of a discipline that is devoted precisely to the representation of entities as they exist in reality

is transformed in footnote 15 of Merrill (2010) into:

“good modeling” *must* be based on a metaphysical realism that embraces universals (emphases added).¹⁹

Under (2), Merrill asserts at various points – on the basis of nothing in our writings – that we claim that studying and embracing our alleged philosophical theory of science and of scientific language is necessary to the proper conduct of science. Some of the users of the realist methodology do indeed concern themselves with such philosophical matters. Some, indeed are former students of philosophy who employ the realist methodology in their work – even though they embrace nominalist positions – *because they see it as bringing practical benefits*. Most, however, do not concern themselves with philosophy at all. And quite rightly so. For we are, like Merrill himself, entirely convinced that no theory of science of the sort produced by philosophers could be necessary to realizing the tasks of science itself.

Under (3), we shall recognize below that, while some of the general terms used in scientific language are to be recognized for ontological purposes as designating types or universals, it is, even in the realm of science (because of the existence of scientific error), not possible to embrace any one-one correspondence between such general terms and corresponding universals or types, and this for rather obvious reasons. We are thus taken aback by Merrill’s assumption that we hold a referential view even in relation to the terms of *natural language*. In his discussion of the two example sentences: ‘John loves Mary’ and ‘John loves pizza’ in (2010), Merrill asserts that, because of the Referentialist Assumption, these sentences are seen by Smith and Ceusters

as being syntactically identical, and so we are urged to conclude that the general term ‘pizza’ must denote some thing (as the individual term ‘Mary’ does) – but not a particular thing ... a universal. (Merrill, 2010, p. 91.)

For it would of course be the height of naivety to apply anything like the Referentialist Assumption to sentences of this sort. Indeed we have argued *ad nauseam* against the drawing of ontological conclusions from the mere surface syntactic features of language. Smith (2005) describes how we see many of the most influential figures of 20th-century analytic philosophy, from Wittgenstein and Carnap to Lewis and Armstrong, as having been

¹⁹Similar creative misquotation is to be found, for example, in (2010, footnote 17), where Merrill asserts that our discussion of certain inadequacies of description logic (Ceusters et al., 2003) ‘attributes any problems (with such logics) to a failure to take seriously the existence and role of universals.’ In fact, however, universals play a role in the mentioned paper only in our discussion of errors of one specific type, namely those which arise through the confusion (familiar under the label ‘*is_a* overloading’) of the relations of instantiation and subsumption.

affected by the erroneous (indeed absurd) assumption that it is possible to infer the ontological structure of reality from the logico-syntactic structure of one specific language.

Further problems turn on the fact that Merrill evinces little first-hand acquaintance either with those practical purposes of ontology development which, on our view, ontologies are primarily created to support, or with the ways the realist methodology is actually being used to solve problems of ontology coordination. The word ‘integration’ appears nowhere in his essay, and neither does any reference to the signature paper – “The OBO Foundry: coordinated evolution of ontologies to support biomedical data integration” (Smith et al., 2007) – in which the application of the realist methodology is described. This apparent ignorance of our actual intentions – and of the reasons for the successes of the realist methodology – give his critique the flavor of one who would harangue soldiers marching into battle on grounds of bad taste in the design of their uniforms.

3. The concept orientation

3.1. International standard bad philosophy

We have alluded already to our agreement with Merrill in the view that no theory of science (or, *a fortiori* of metaphysics) of the sort produced by philosophers could be necessary to realizing the tasks of science itself. We thus share with him the view that it would be inappropriate for philosophers of science – or metaphysicians of whatever stripe – to attempt to interfere with how scientists do their job.

For this very reason, however, we have become acutely conscious in our work with various communities of scientists of the degree to which scientific conduct *is* being interfered with philosophically on another plane as a result of the increasing importance to science of computational artifacts. This interference comes not from the side of philosophy itself, however, but rather from information and computer science (Smith, 2004). Much of the polemical work we have published in recent years has been addressed to the task of counteracting this interference as it emanates especially from disciplines such as knowledge engineering and conceptual modeling – disciplines which exert a strong impact especially on the ontology field and thus, indirectly, on science.

We have identified in this connection a collection of views that we have labeled ‘International Standard Bad Philosophy’ (Smith, 2005), views first clearly identifiable in the thinking of one Eugen Wüster, a businessman, Esperanto enthusiast and nominalist, and also the founder in 1951 of the Technical Committee (TC37) for terminology standardization of the International Standards Organization (ISO). Wüster’s views on the proper method of terminology development enjoyed a quite astonishing influence, which persists in many terminology standardization efforts even today, as can be seen from many passages in contemporary ISO standards documents, such as the following (from 1999) taken over almost verbatim from earlier writings of Wüster:

an object (for purposes of terminology work) is defined as anything perceived or conceived. Some objects, concrete objects such as a machine, a diamond, or a river, shall be considered material; other objects are to be considered immaterial or abstract, such as each manifestation of financial planning, gravity, flowability or a conversion ratio; still others are to be considered purely imagined, for example, a unicorn, a philosopher’s stone or a literary character.

... In the course of producing a terminology, philosophical discussions on whether an object actually exists in reality ... are to be avoided. Objects are assumed to exist and attention is to be focused on how one deals with objects for the purposes of communication. (ISO, 1999, emphasis added.)

If, however, terminology standards are constructed in such a way that real objects such as rivers are placed on the same level as imagined objects such as unicorns, then it is unlikely that the terminologies that result will be able to support the current needs of, for example, biological science.

The most pervasive influence of International Standard Bad Philosophy is via the doctrine according to which terms in ontologies should be seen as referring, in some sense, to *concepts*, as captured, for example, in the ISO definition of a terminology as a ‘set of terms representing the system of concepts of a particular subject field’ (ISO, 1990).

It is because this doctrine has given rise, and continues to give rise, to multiple false steps in the discipline of ontology – false steps that we see being repeated over and over again in every new area in which ontology technology is applied – that we have devoted so much effort to developing and disseminating an alternative, more scientifically coherent, view as to how the terms in ontologies should properly be understood, in the hope that such false steps can be avoided in the future.

At one point in “Ontological Realism” Merrill remarks of the realist approach that, while it ‘may be looked upon favorably ... by medical informaticists who lack familiarity with alternative approaches and who – for a time at least – may be enticed into going along for the ride, empirical scientists (will) find it much more difficult’. Interestingly, however, it is precisely among medical informaticians and computer scientists that we find the most visceral resistance to the realist approach. Empirical scientists, in contrast, have been supportive of our efforts from the beginning (and the reader is invited to note the large number of bench biologists in the list provided at the end of Section 1.10).

Why should this be so? Why, more precisely, should so many informaticians and computer scientists (and terminologists) remain so faithful to the concept orientation and, more generally, to one or other subjectivist or relativist view, conceiving ontologies as representations not of some independent reality but rather of mere *views* or *perspectives* or *descriptions* or ‘collective hunches’ (Smith, 2004)? Why, on the other hand, should so many bench biologists be so open to the realist alternative?

Part of the answer lies, we believe, in the fact that computer scientists – unlike most biologists – receive training in cognitive psychology, which leads encourages them to have strong feelings about what they see as the constructed nature of much of human belief. Another part has to do with the existence of incentives within the world of information technology which support the creation of new intellectual resources rather than the refinement and reuse of those which already exist. For empirical biologists, on the other hand, incentives often point in the opposite direction, which means toward finding ways to ensure that past, present and future data can be effectively shared.

3.2. The National Cancer Institute Thesaurus

One of the first applications of the methodology of ontological realism was to the critical analysis of the National Cancer Institute Thesaurus (NCIT) (Ceusters et al., 2005), a component of the UMLS Metathesaurus collection of biomedical source vocabularies, in which we identified a series of embarrassing errors of definition, classification and logic. The NCI commissioned an overhaul of its Thesaurus in response to our criticisms, though the organization contracted to make the needed changes thereby succeeded, in some respects at least, in making matters worse.

Many of the errors of the NCIT, then and now, grow out of confusions surrounding the term ‘concept’ and its cognates, as for example in the use-mention confusions present in NCIT definitions such as:

Conceptual entity =_{def.} An organizational header for concepts representing mostly abstract entities.

Event occurrence =_{def.} An indication or description that something has occurred.

These formulations (taken from the version of NCIT current in June 2010) are not only logically nonsensical (they are comparable to, for example, ‘Swimming is healthy and has two vowels’); they are also practically useless for anyone who might want to understand how, precisely, the respective terms are intended to be used by the authors of the NCIT.

The confusions manifest themselves also in the circular *is_a* relations present in the NCIT, as for example in:

Entity *is_a* Conceptual Entity,

an assertion logically comparable to: ‘apple *is_a* green apple’.

As we have repeatedly urged, the reason why there are so many errors associated with NCIT’s use of ‘Concept’ – just as there are so many parallel errors in other parts of the UMLS Metathesaurus – is because the authors of the NCIT do not understand what they are referring to when they use the word ‘Concept’. One prime indication of this lack of understanding is the number of occasions on which items classified by the NCIT under ‘Concept’, ‘Conceptual Entity’ or cognate terms are misclassified or inconsistently defined. Neither ‘Concept’ nor ‘NCI Administrative Entity’ is classified as a Conceptual Entity in NCIT, for example, and this even though the latter is *defined* by the NCIT as ‘Conceptual entities (*sic*) required by NCI operations and systems’.

Some of the nicest examples of ‘Conceptual Entity’ terms in NCIT are found among the children of ‘Geographic Area’, which include ‘Alabama’, ‘France’ and ‘door’.²⁰ One troubling issue here – troubling because it suggests that the authors of the NCIT have an uncertain understanding, not merely of geography, but also of the basic rules of logic – is that Alabama is asserted to be a *subclass* of US State, just as France is asserted to be a *subclass* of Country (and just as the Burgundy wine region is asserted, in Noy and McGuinness (2001), to be a subclass of France). States, countries and wine regions, however, are not classes on any of the normal understandings of ‘class’; and thus also they are not subclasses of other classes.

3.3. The future of the concept orientation

Of course the argument according to which adoption of the concept orientation in ontology and terminology development thus far has been associated in this way with certain characteristic errors is not as yet an argument to the effect that the concept orientation must be abandoned in its entirety. It shows only that we need a more carefully formulated, and technically more sophisticated, and consistently disseminated, account of what concepts are, perhaps of the sort that is attempted in Merrill (2009).

Currently, however, matters do not bode well for the concept-based methodology. In the most recent (2010) release of SNOMED CT, for example, hitherto one of the most

²⁰<http://nciterns.nci.nih.gov/ncitbrowser/ConceptReport.jsp?dictionary=NCI%20Thesaurus&code=C48950>, last accessed August 20, 2010.

enthusiastic implementations of the concept orientation, the included Dictionary provides in its entry for ‘Concept’ the following definition:

Concept: An ambiguous term. Depending on the context, it may refer to:

1. A clinical idea to which a unique ConceptId has been assigned.
2. The ConceptId itself, which is the key of the Concepts Table (in this case it is less ambiguous to use the term ‘concept code’).
3. The real-world referent(s) of the ConceptId, that is, the class of entities in reality which the ConceptId represents (in this case it is less ambiguous to use the term ‘meaning’ or ‘code meaning’).²¹

Simultaneously, figures as influential in the medical informatics and international standards (ISO) communities as Chris Chute and Harold Solbrig are now recognizing that there is something very wrong with this same concept-based methodology, and they have documented their concerns in a paper entitled “Concepts, Modeling and Confusion” (Solbrig & Chute, 2009). Commenting on this paper Solbrig reports how:

When I first started working on the (National Center for Biomedical Ontology) project I didn’t fully buy in to the realist approach. The process of resolving your critique of the NCI Thesaurus, however, convinced me that from a purely pragmatic perspective the approach (mostly) worked. Since then, I have continued to apply some of the basic organizational principles and have been pleasantly surprised at how useful they have been in defining, organizing and classifying all sorts of knowledge resources. Somewhere along the way it just became intuitive and obvious – science is about describing reality, and the primary point of agreement has to be on the things being described. I have to admit that I still don’t agree with some of the techniques that have been used to publicize this approach, but it is obvious, however, that what you have been doing is working.²²

4. Merrillian alternatives to the realist methodology

4.1. Carnap

Merrill does not himself advance a strategy for ontology coordination. As far as one can tell from his (Merrill, 2009) and (Merrill, 2010), however, for such a strategy to be capable of receiving Merrill’s support it would have to be centered on the use of FOL. The most ambitious such strategies involve the translation of scientific content into the language of FOL along lines first attempted by Carnap in the service of what was in his day referred to as the Unity of Science Movement (Morris, 1960). In *The Logical Structure of the World* (Carnap, 1928), he offers a methodology for translating all of science into one single ontology based on a doctrine called ‘resemblance nominalism’. The approach uses Carnap’s own dialect of the language of FOL, which differs in two respects from that of received FOL, first in allowing terms to represent what he calls ‘elementary experiences’, second in allowing only one single primitive dyadic predicate ‘M’, which is satisfied if and only if two particulars ‘match’ each other. The result set standards of logical rigor and of syntactic constraint in the service of the integration of the content of scientific theories which remain unsurpassed. But Carnap’s method nonetheless failed, not least because of what Carnap’s fellow nominalist Goodman called the ‘disastrous’ problem of ‘imperfect community’

²¹Previously, SNOMED CT had defined ‘Concept’ as: ‘a unique unit of thought’. At the same time it defined ‘Disorder’ as: ‘a concept in which there is an explicit or implicit pathological process causing a state of disease which tends to exist for a significant length of time under ordinary circumstances.’ From this it can be inferred that some *units of thought* contain *pathological processes* causing *states of disease*.

²²<http://www.bioontology.org/node/540>, last accessed June 30, 2010.

(Goodman, 1951), a problem turning on the fact that simple examples can be constructed to show that given groups of particulars may resemble each other yet fail to share any property in common. Moreover, because two putatively distinct universals may happen to have exactly the same instances. Carnap's method of constructing natural classes on resemblance-nominalistic principles would then incorrectly determine only one class for what intuitively seem to be two universals (thus: two respects in which the same things resemble one another).

Among the community of those currently attempting to construct ontologies in support of integration of scientific data, therefore, Carnap's method has no actual users. The general approach of translation of scientific content into the language of FOL was however revived in the work of McCarthy and Hayes (1969) and other pioneers in the field of ontology as a discipline allied to the study of artificial intelligence and to what is called 'knowledge representation' (Gruber, 1992). Such an approach, when freed from the constraints of Carnap's resemblance nominalism, has a number of points in its favor, not least being the fact that it combines the rigor of predicate logic with a high degree of expressivity. Unfortunately, however, this very expressivity comes at a price. This is because unrestricted FOL, even on the received view of what its terms are allowed to represent, is still so highly expressive that for many sentences it can be used to formulate a plurality of regimentations of a sort that defeats the goals of coordinated ontology development (see Section 4.3.2). This problem is only partially resolved if one uses a language such as the Web Ontology Language (OWL), which, although logically less expressive than FOL, still contains no in-built constraints on the sorts of predicates that can be incorporated into the ontologies built in its terms. The FOL-based approach is still alive in the SUMO initiative (see Section 6). Significantly, however, as we shall see, SUMO has failed to capture the enthusiasm of biologist users who would be willing to invest the efforts necessary to use received FOL à la SUMO as a basis for formulating the content of biological science and of using the result to annotate actual data.

4.2. Woodger

An approach which, in terms of constraints on expressivity falls somewhere between Carnap and SUMO was put forward by J.H. Woodger, another member of the Unity of Science Movement, and also translator of Tarski and Reader in Biology at the Middlesex Hospital Medical School in the University of London. In his remarkable *Axiomatic Method in Biology* (Woodger, 1937), Woodger provides a translation into predicate logic of major portions of the biology of his day, thereby anticipating in terms of formal rigor, generality of scope, and scientific coherence much later achievements of logically-informed biomedical informatics such as, for instance, the GALEN project (Rector & Nowlan, 1994).

Sadly, however, Woodger's initiative, too, must be judged a failure, and this for a number of reasons. First it was some 50 years ahead of its time, since the potential utility of the sort of formalized representation attempted by Woodger became manifest only with the widespread use of computation in support of scientific research. Second, Woodger's axiomatization falls short from the point of view of modular organization, so that there is missing any distinction between formal-ontological (top-level, organizing) portions of the theory and domain-specific portions corresponding to the separate biological disciplines.

All terms used in the formulae of Woodger's theory are defined in terms of the small set of primitives listed in Table 1. This provides a promising approach to the creation of the sort of constraint on expressivity that is needed if the goals of integration are to be achieved. But at the same time the various domain-specific portions of Woodger's theory – for example, its treatments of Mendelian genetics and of embryology – are so intricately embrangled with each other in the formalization that, were one portion of the theory to be rejected because of

empirically-based advances in the relevant parts of science, then the entire theory would have to be rejected also. For the same reason, too, Woodger's approach does not lend itself readily to the sort of division of labor which would allow distinct components of the theory – for example, in cell biology or in evolutionary systematics – to be developed in a dedicated fashion by experts in the corresponding disciplines.

All of which brings us to what is from our present perspective the principal problem with Woodger's approach: the absence of modularity – or of what we could now call 'normalization' – brings not only obstacles to the theory's being able to keep pace with scientific advance; it implies also that – as Fig. 3 makes clear – his theoretical contribution, as expressed in page after page of logical formulae, is practically impenetrable to all but a very small minority of specialists in mathematical logic. Our experience working with ontologists and scientists in biological and similarly complex domains has taught us, however, that there is an essential trade-off between logical complexity on the one hand and biological usability and revisability on the other. There was then, and is now, no way in which Woodger's contribution could have been useful to biomedical researchers. For given the scalability problems of the biomedical ontology integration task, ontology resources will require at every stage significant contributions from multiple disciplinary groups of biologists who are in a position to ensure that these resources are properly maintained and properly used. Ontologies will receive the support they need from biologists in this way, however, only if the latter are able both to *understand their contents* and have confidence that they will evolve in such a way as to keep pace with scientific advance. Ontologies which do not capture the relevant audiences of human users, even if they achieve very high standards of technical rigor, will for scientific purposes be as worthless as, for example, telephone networks meeting the highest of technical standards but with no actual subscribers.

4.3. How would Merrill approach the task of ontology coordination?

Merrill's purpose in "Ontological Realism" is a negative one. It is to demonstrate that the realist methodology, while it contains several elements of which he approves, also contains other elements – centered on uses of the word 'universal' – that are subject, as he sees it, to serious flaws and therefore ought to be abandoned.

We do not at all rule out that there might be ingredients in our methodology that are inessential to its proper functioning and perhaps even detrimental in this or that way. We are dealing, after all, with a large-scale effort in scientific coordination, where multiple path dependencies will play a necessary role. But it is not at all clear that Merrill himself has succeeded in identifying any such detrimental elements; and even if he had, we would be reluctant to make any attempt to untangle them from the whole without good evidence of what might be the consequences of such an attempt.

When we examine the content of Merrill's critique, however, we find too little that is of substance to justify such a change. For, when we leave aside his recommendations concerning logic and semantics – with many of which, were they only clearly specified, we would almost certainly agree – this critique amounts to a rather peculiar argument, resting in no small part on a series of misquotation of our writings, which we might characterize in a preliminary form as follows:

- Smith and Ceusters, in writing about their methodology, occasionally refer positively to Aristotle, and to David Armstrong, and, like them, they use the word 'universal' to formulate their realist views (true).
- Some of Aristotle's and Armstrong's ideas are inconsistent with empirical science (true).

- Therefore, Smith and Ceusters in using the word ‘universal’ when describing the ontological realist methodology cannot possibly be doing anything which helps empirical scientists to do their work (false).

4.3.1. Adherence to the principles of logic and semantics—In his (2007) Merrill described his work (on the GlaxoSmithKline Babylon Knowledge Explorer) as a kind of scissors-and-paste engineering. If, for example, the GO, or the WHO Drug Dictionary, are found to be flawed, then our response, he says, ‘cannot be to devote time and effort to repairing such flaws in a systematic manner. Instead, it is to work with what is available or ... to make what is available work’. At this stage in the development of his thinking, therefore, Merrill seemed to hold out no hope for the realization of the goal of ontology coordination that is at the center of our work. If this is his position today, and if he has arguments for this position, then we would assume that he would find it most sensible to criticize our methodology on the basis of these arguments rather than on the basis of incidental features of our writings still bearing traces of a philosophical etiology. We speculate therefore that he has moved on from the total pessimism of 2007, and accordingly focus on those passages in his writings which can be interpreted as allowing for the possibility of some sort of ontology coordination strategy analogous to our own.

The first such element can be formulated as follows: that to develop ontologies able to meet the needs of biomedical research, authors need to ‘understand and employ the principles of formal logic, semantics, and the philosophy of language’. We will, thereby, Merrill says, ‘avoid the confusions and errors that Smith and Ceusters have quite rightly criticized in a number of flawed approaches to ontologies in science’ (Merrill, 2010, p. 105).

Unfortunately, however, this is not so. Indeed, as Merrill himself is fully aware, it is not even clear that there are commonly accepted ‘principles of formal logic, semantics and the philosophy of language’. As he himself expresses it (personal communication): ‘There are a number of ways of approaching the semantics of sentences, of terms and of predicates. Many of these ways are incompatible with one another, and each has certain advantages, disadvantages and challenges’.

The proposal as stated is also marked by a certain naivety as concerns the work which must be done if those engaged in ontology development in the service of science are indeed to be brought to the point where they *are truly able to avoid* confusions and errors of the sorts we have identified. For we have ample evidence that even those schooled in the practical application of the disciplines of logic and semantics may fail to recognize the need for ontologies that enjoy, for example, the feature of mutual consistency; some, indeed, are creating ontology-like artifacts which are unashamedly not internally consistent even with themselves (Lenat, 1995).

4.3.2. The focus on predicates rather than on general terms—In a number of places Merrill recommends versions of the ‘principle of tolerance’ articulated by the later Carnap as follows: ‘Let us grant to those who work in any special field of investigation the freedom to use any form of expression which seems useful to them, ... and tolerant in permitting linguistic forms’ (Carnap, 1950). While such a principle is of course perfectly acceptable in the context of hypothesis-driven experimental science, it would be the kiss of death in the context of ontology. For as we have claimed already above, and as we argue in detail in Section 6, ontology-based integration of data in a complex and heterogeneous domain like that of the life sciences is in practice unachievable except through the application of constraints on what can be said within the framework of the ontologies created. Merrill’s endorsement of the tolerance principle will thus be seen to mark yet

another worrying element of naivety on Merrill's part when it comes to addressing the needs of real-world ontological development.

Interestingly in both Merrill (2009) and (2010) seems to offer arguments not obviously compatible with the spirit of the principle of tolerance in favor of the merits of a regimentation of the content of ontologies and terminologies that would be based, not on general terms, as is standardly the case, but rather on predicates. To see what this would mean, consider the sentence:

(H) Lipitor is an HMG-CoA reductase inhibitor.

In ontologies modeled after the GO this sentence would be regimented via an assertion linking two nouns, for example as follows:

(I) Lipitor *has_function* HMG-CoA reductase inhibitor.

On Merrill's proposal, in contrast, it would be rendered as a universally quantified FOL statement linking two predicates:

instantiates_Lipitor

instantiates_HMG-CoA reductase inhibitor

to the effect that everything which satisfies the first predicate satisfies also the second. In symbols:

(J) $(\forall x)(\text{instantiates_Lipitor}(x) \rightarrow \text{instantiates_HMG-CoA reductase inhibitor}(x))$.

There is now one obvious reason why all successful ontology and terminology ventures in support of science thus far have preferred the first, general term-based, approach. Consider a sentence such as

Simvastatin activates *the protein kinase Akt* and promotes *angiogenesis* in *normocholesterolemic animals*.

Here the number of general terms that can be identified is rather limited – relevant candidates have been italicized. In the case of predicates, in contrast, because the latter can be combined with each other to form logically more complex predicates in multiple arbitrary ways, there will be, for any sentence of reasonable complexity, indefinitely many logically acceptable predicates that can be identified within it. In the mentioned sentence, for example, we can identify predicates such as:

activates the protein kinase Akt

activates the protein kinase Akt and promotes angiogenesis

promotes angiogenesis

promotes angiogenesis in normocholesterolemic animals

activates the protein kinase Akt and promotes angiogenesis in normocholesterolemic animals

is promoted by Simvastatin

is activated by Simvastatin

is activated by something

promotes something

promotes something in normocholesterolemic animals

and many more.

Embracing a predicate-focused approach to the logical regimentation of scientific content thus gives rise to the same increased likelihood that ontological representations will fork that we have identified already in our discussion of the concept orientation above. The predicate-based approach does indeed allow biologists to say what they want. The problem is that they can say *many more things also*, and this approach provides no guidance as to how the needed selection is to be made in a reliably coordinated fashion by dispersed groups of ontology authors.

And worse. Given that predicates, like concepts, and like classes in intension, can be combined unrestrictedly by means of Boolean operators such as ‘and’ and ‘not’, there is nothing to rule out the appearance in predicate-based ontologies of absurd combinations such as ‘instantiates liver and tree’. Hogan (2009) points to an example of such an absurdity in the then current version of SNOMED-CT, where albumin-bound paclitaxel is referred to as a subtype of both albumin and paclitaxel, where of course, no molecule is both an instance of albumin and an instance of paclitaxel at the same time.

4.3.3. Privileging FOL—Merrill advocates in many places the use of FOL (or of the related Common Logic family of first-order logics), just as he recommends authors such as Cocchiarella (2003), Zalta (1983) and Lenat (1995), whose work in logic, philosophy or computer engineering is rooted in the use of FOL or of the second order logics associated with FOL.²³ In one passage addressing the relations between FOL and other kinds of logic Merrill reveals particularly clearly how little he has familiarized himself with the actual practices of contemporary biomedical ontology, and specifically his apparent ignorance as concerns the role and nature of the different sorts of logic that are employed therein:

The Referentialist Assumption [which Merrill sees as being adopted by Smith and Ceusters] makes more sense if it is adopted against the background of a *term* logic (such as Aristotle’s syllogistic, the logic of Leibniz or Boole, or Description Logic) rather than a *predicate* logic (such as modern first-order predicate logic) [1]. Term logic, for good reason, has been referred to by Peter Simons as “logic lite”; and its weaknesses are well known (among them, lack of expressive and inferential power) [2]. Its reintroduction in contemporary times as Description Logic [3] has been an attempt to provide a simplified formal basis for automated reasoning, but its flaws are proving to be too high a price to pay in many applications [4], and so alternatives are being sought – among them, Common Logic [5]. In other contexts, when one is willing to pay the price in terms of computational resources and performance, standard first-order logic (or something even stronger, as in the case of Cyc which adds some second-order extensions ...) provides a much more satisfactory framework for knowledge representation and reasoning [6]. (Merrill, 2010, footnote 17.)

Ad [1]: As will become clear in Section 5.4 and as is documented at length in Smith (2005),²⁴ our highly constrained version of what Merrill calls the “Referentialist Assumption” is anchored entirely within FOL.

Ad [2]: Those who have familiarity with the role of logic in major ontology development projects will know that it is vital for some purposes to have at one’s disposal a basic logical resource that has the (‘weak’) expressive power and support for inferencing that is needed for publishing and editing of ontologies; this is proved not least by the tremendous success of the OBO format, and of the OBO-Edit software resource,²⁵ to which Merrill nowhere

²³Oddly, the first two of these authors defend versions of metaphysical realism considerably more extreme than the version of this position that Merrill imputes to Ceusters and Smith – and criticizes so vehemently.

²⁴A paper Merrill refers to in exactly this connection in Merrill (2010, footnote 9), but seems not to have read.

²⁵<http://oboedit.org/>, last accessed June 30, 2010.

refers, even though they continue to serve as the principal pipeline through which high quality ontology-annotated data enter into the public domain.

Ad [3]: This is an egregious error, since Description Logics – in the plural – have nothing whatever to do with term logic but are rather a (family of fragments of) FOL which have the very same FOL semantics (modulo constrained expressivity) and in which predicates – including the relational predicates absent from term logics – play the very same (Merrill-approved) role.

Ad [4]: The Description Logics used within the biomedical ontology development community are primarily the OWL (Web Ontology Language) with the profiles OWL-EL and OWL-QL according to the new OWL-2 specification.²⁶ SNOMED CT (roughly) uses the less expressive OWL-EL variant. Users of Description Logics are aware of the many issues which flow from the constraints on expressivity that are imposed for the sake of certain vital computational benefits, which include the facility, when working in OWL 2.0 and in certain other Description Logics, to check successive ontology drafts for consistency in ways guaranteeing a response that is in almost all cases close to immediate, and to import and export ontology content in flexible ways (Courtot et al., 2011). Certainly it is true that these constraints on expressivity have often led to embarrassingly trivial work.²⁷ Problems arise also in virtue of the often seemingly willfully confusing choices of technical metaterminology by the authors of OWL, for example, using ‘property’ for what in other circles is called ‘relation’. Because of these factors Smith and Ceusters initially belonged to the camp of skeptics as concerns the use of OWL in scientific contexts. Largely as a result of the efforts of those working within the OBO Foundry community, however, an impressive and ever-increasing body of scientifically valuable content is now available on the web using OWL as native development format.

Ad [5]: Again, because of his lack of familiarity with the body of work that he sees fit to criticize, it is in fact, as concerns the life sciences, precisely within the community of users of the OBO format that experiments in the use of Common Logic as a resource to supplement the expressivity of weaker logics are being made,²⁸ and in ways exploiting aspects of the realist methodology which Merrill assails.

Ad [6]: We will return below to consider the merits of Cyc, or of anything like Cyc, for purposes of scientific data integration. There are many reasons why, after some \$100 million of investments in its development, there are still no documented successes in this regard on the part of Cyc. One reason is of course that Cyc was built for a quite different purpose. Another reason, as concerns biology at least, is its content, of which we here provide just one sample, taken at random²⁹:

`BiologicalReproductionEvent =def A specialization of BiologicalProductionEvent. Each instance of BiologicalReproductionEvent is an event in which one or more instances of BiologicalLivingObject (q.v.) (related to the event by parentActors) produce at least one new instance of BiologicalLivingObject (related to the event by offspringActors), generally of the same kind as the parents.`

`ConceivingSomething_BiologicalReproductionEvent =def a collection of events; a subcollection of BiologicalReproductionEvent. In each ConceivingSomething_BiologicalReproductionEvent, someone becomes pregnant.`

²⁶<http://www.w3.org/2007/OWL/wiki/Profiles>, last accessed August 10, 2010.

²⁷Some examples of OWL ontologies which seem to have been created without the help of the realist methodology can be examined here: <http://www.schemaweb.info/>, last accessed June 30, 2010.

²⁸<http://github.com/cmungall/bio-clif>, last accessed August 10, 2010.

²⁹<http://sw.opencyc.org/concept/Mx4rU7IAg7HiEdmAhAACs6hRjg>, last accessed August 10, 2010.

The immaculate conception =_{def.} The ConceivingSomething_BiologicalReproductionEvent in which Mary_MotherOfJesus was conceived. Catholic dogma holds that Mary (unlike Jesus) was conceived by conventional biological means, but that GodOfAbrahamIsaacAndJacob interceded at the time of her conception to keep her free from the stain of original sin, or ‘immaculate’.

It is a poignant expression of Merrill’s naivety when he recommends the Cyc resource – in the context of a critique of a successful, practical strategy for addressing problems of cross-disciplinary data integration in the biomedical domain – as providing a ‘more satisfactory framework for knowledge representation and reasoning’.

We have distinguished three requirements which Merrill, it seems, would recommend for any ontology development project: adherence to the principles of logic and semantics, the focus on predicates, and the privileging of FOL.

We believe that observance of any of these requirements would, each for different reasons, guarantee failure for any strategy for ontology coordination constructed in its terms: the first because it is so woefully underdetermined, the second because it will guarantee forking, the third because, at this stage in the development of logical and ontology technology, at least, FOL-based initiatives can be made useful in the work of biomedical ontology development only if they are employed in tandem with (or as precursors to) the development of simpler logical resources with certain needed computational benefits.

5. Merrill’s misunderstandings of the realist methodology

5.1. Armstrong

It was above all the philosopher David Armstrong who pioneered the thesis according to which the study of universals might be of value in the defense of scientific realism (Armstrong, 1978). And just as we find the broad frame of Aristotelian realism congenial, so we find much that is of value in Armstrong’s writings. At the same time, however, we have devoted considerable effort to the criticism of Armstrong’s thinking (Smith, 2005), and we are thus puzzled by Merrill’s assertions – which play a crucial role at multiple points in his exposition – to the effect that ‘the realism of Smith and Ceusters is explicitly modeled on Armstrong’s metaphysics’, and that it is Armstrong’s metaphysics which ‘Smith-Ceusters takes to be fundamental to the theory of universals being embraced’ (Merrill, 2010, p. 88). This is, we are sorry to say, just not true – and the assumption that it is true, unfortunately, renders much of Merrill’s critique of our position irksomely irrelevant.

Specifically, we depart from Armstrong’s views in at least the following crucial respects (Smith, 2005; Neuhaus et al., 2004), including:

1. the central role he awards to the ontological category of *states of affairs* or *facts*, which he views, oddly, as constituting the ultimate simples in the universe,
2. his view of universals and particulars as (quasi-epistemological) dependent parts or aspects of states of affairs,
3. his reliance upon a mythical ‘future perfected state of science’ in which, as he sees it, his own formal-ontological proposals will be finally realized,
4. his concomitant failure to address realistic examples taken from really existing sciences such as physics or biology,

5. his assumption that all universals are properties or attributes (Armstrong, 2008), and thus entities corresponding – albeit not via any one-to-one mapping – to predicates,
6. his concomitant rejection of what Aristotle identified as universals in the category of substance, such as *molecule*, *cell*, *organism* or *planet*,
7. his unquestioning assumption – embraced also by Merrill – of the serviceability of received predicate logic as a template for creating ontologies. In Smith (2005) we refer to this assumption under the label ‘fantology’, in light of the fact that there is, for Armstrong as for Merrill, no role for general terms in the properly regimented language of science, but rather only for predicates (‘F’) and singular terms (‘a’). This reflects the thesis at the center of the received interpretation of FOL, according to which *all generality lies in the predicate*, and never in the subject.

To see how Merrill’s criticisms of our ontological views fall wide of the mark because he imputes to us Armstrongian positions we do not hold, consider, for example, (3), above. Where Armstrong can in all seriousness hold that to establish what universals there are we need to appeal to the future perfected state of what he calls ‘total science’ (Armstrong, 1989, p. 87), we ourselves are interested precisely in really existing scientific theories, and in the associated really existing ontologies, which in normal circumstances are not associated with any claim to completeness. Really existing scientific theories are marked, rather, by messy and inconvenient processes of change and of correction of error, including ontological error, and our formulation of the realist methodology is designed precisely to do justice to this fact (Ceusters, 2009). Where Armstrong’s views are put forward as *philosophical doctrines*, ontological realism is a *practical methodology*. In order to sustain his attack on ontological realism on grounds derived from flaws he finds in philosophical doctrines defended by Armstrong, therefore, Merrill is forced into contortions of positively Ptolemaic proportions.

Certainly we share Armstrong’s recognition of the need for a *sparse theory of universals* – as contrasted with those theories which allow representations of universals/properties/intensions/concepts to be constructed in combinatorial fashion. Armstrong himself formulates the sparse view as follows: ‘Given a *predicate*, there may be none, one or many universals in virtue of which the predicate *applies*’ (Armstrong, 1978, emphasis added). For us the sparse theory is a view to the effect that for each *scientific general term*, there may be none, one or many universals to which the general term *refers* (Smith, 2006a). We, like Armstrong, hold the sparse theory of universals because of our conviction that the question as to which universals exist in reality is a matter for scientists, not for ontologists, logicians or linguists, to determine (Grenon & Smith, 2004). Unlike Armstrong, however, for whom what matters is the future ‘total science’ when scientists will exist in a state of epistemological perfection, we acknowledge that it is impossible to read off from any given scientific theory what universals exist in reality for simple epistemological reasons – turning on the fact that the theory in question may rest on error.

5.2. The non-smoker

From the ontological realist perspective, that a specific universal exists is never a matter of what can be discovered by logical means alone, but always only through application of the scientific method.

In particular, therefore, we reject the thesis according to which, from the fact that F is a universal, we could infer that non-F is a universal, where ‘non-F’ is defined as follows:

³⁰See, for example, Merrill (2010, footnote 13).

(K) x instantiates non-F =_{def.} it is not the case that (x instantiates F).

Indeed we go further and argue that:

(L) If 'F' designates a universal then 'non-F' (in the sense defined by (K)) does not designate a universal.

(L) implies, in particular, that if 'smoker' designates a universal, then 'non-smoker' (in the sense of (K)) does not designate a universal.

Here Merrill sees trouble for our position, in light of the fact that assertions such as:

(M) Non-smokers are less susceptible to cardio-pulmonary diseases than are smokers, might very well be supported by empirical evidence. From this, he infers, it follows that ontological realists might potentially be in a position where they would have to reject empirical evidence because it would contradict some favored metaphysico-logical principle. If he were right in this, then (L) would of course need to be sacrificed, and Merrill, because he would have finally discovered an actual error in our work, would have scored a valuable point.

Unfortunately, however, in making this charge Merrill confuses what are standardly called 'internal' and 'external' negations, and thus himself commits an error of logic.³¹ This is because 'non-smoker', as it occurs in assertions such as (M), utilizes only the *internal* negation expressed by *human who does not smoke*, not the external (which is to say logical or Boolean) negation conveyed by: *entity of which it is not the case that it smokes*. A cardinal number, or a glass of water, is a non-smoker in the latter sense, which is the sense captured by (K); not however in the former.

Our assertion (L), now, has no implications at all for terms (such as 'odorless', 'colorless', 'invisible', 'unfriendly' and so on) involving mere internal negation, since the sparse theory of universals of which (L) is one expression pertains only to the question whether representations of universals can be composed through application of logical constants such as 'and' or 'not'. And we are confident that every assertion analogous to (M) in which 'non-smoker' or any similar term would truly be to be interpreted in the externally (i.e., logically) negated sense, will be found to be clearly false on the basis of simple inspection. Thus it was, the last time we checked, not the case that cardinal numbers are less susceptible to cardio-pulmonary diseases than are smokers.

In Ceusters (2007) we set out the recommended realist treatment of negative assertions. Both 'smoker' and 'non-smoker', if included in an ontology conformant to the principles presented above, would need to be included in the corresponding inferred hierarchy on the basis of definitions along roughly the following lines:

$$\text{smoker}(x) =_{\text{def.}} \text{instantiates}(x, \text{human being}) \ \& \ \exists y((\text{instantiates}(y, \text{act of smoking}) \ \& \text{participates}(x, y)))$$

$$\text{non-smoker}(x) =_{\text{def.}} \text{instantiates}(x, \text{human being}) \ \& \ \neg \exists y((\text{instantiates}(y, \text{act of smoking}) \ \& \text{participates}(x, y)))$$

employing universals *human* and *act of smoking*. We say 'roughly' because a full account would need to specify the thresholds for when somebody would count as belonging to one or other collection, for example in the case of humans who *recently gave up smoking*, or who *smoke occasionally*. (M) on this account would then amount to an assertion relating *acts of*

³¹See Slater (1979). In Ceusters et al. (2005) we show how this same logical error is committed also by the curators of the NCI Thesaurus.

smoking repeated above certain frequencies to elevated risks of cardiopulmonary disease on the part of the corresponding individuals.

5.3. Ostrich nominalism

Consider the sentence:

Teco is a bonobo,

and let us assume that this sentence is true. We can then ask what it is in reality that makes it true (and thus what sorts of things scientists would have to attend to, in reality, in order to verify that it is true). In part, clearly, Teco. But on some (non-nominalist) accounts, there is in addition a second something that contributes to making true the sentence in question, some *feature* or *way of being*, or some *species* or *natural kind* to which Teco belongs, or some *structure* or *pattern of DNA* in Teco's genome in virtue of which it is true that he is a bonobo. The sentence in question then asserts a relation between Teco and this second something.

In "On What There Is" (Quine, 1953), Quine presents an alternative to views of this sort designed to lend support to nominalists, like himself, who have a taste for austere ontologies. For the world to be such that Teco is a bonobo, Quine holds, it must be the case that the world includes some bonobo; but it need not include anything properly referred to by means of a general term such as, say, 'bonobohood' or '*Pan paniscus*'.

From Quine's point of view, 'A subject-predicate sentence is true if and only if the subject satisfies the predicate'. Thus, for example, 'Snow is white' is true if and only if snow is white. Many have been disconcerted by the apparent circularity of this doctrine. Armstrong (1978) gives voice to this puzzlement by coining the term 'ostrich nominalism' as a label for those philosophers 'who refuse to countenance universals but who at the same time see no need for any reductive analyses' of the sort that would replace talk of universals for example with talk of sets or collections of resembling particulars.

For Armstrong, questions like 'what makes it true that Teco is a bonobo?', or more generally, 'what is it for *a* to be an instance of the type *T*?', or 'for *a* to have the attribute *F*?' are *compulsory questions* – questions that all upstanding philosophers are called upon to address (Armstrong, 1980). The ostrich nominalist's response to such questions, however, is to bury his head in the sand – while everyone else in the debate (even the most extreme of nominalists who might appeal, for example, to brute relations of resemblance) thinks that *a*'s being *F* warrants some form of analysis.

In response, the ostrich nominalist might argue that, on his account, the phenomenon of true predication is a basic phenomenon, one not reducible to, or explainable or analyzable in terms of, anything more fundamental. Circularity is, in this sense, both inevitable and harmless. In fact, however, we think that the only reason for treating predication in this way as brute (which is to say: not further analyzable) comes, again, from an overblown fascination with austere ontology (with a taste, as they say, for desert landscapes). Such an ontology is revisionary; and it is adopted by the ostrich nominalist without good reason (and certainly without any reason being *supplied*).

Merrill, too, where others see general terms, professes to see only predicates with no referential force. In the sentence 'Socrates is a man', he writes,

the term 'Socrates' is singular and denotes a particular man while the term 'man' may be taken to be a general term denoting the class of men, the form *Man*, or mankind. Alternatively, in modern first-order logic, 'man' would not be regarded as a term in this sentence, but rather 'is a man' would be regarded as a predicate. And the difference here

is that predicates are not (or certainly need not be) viewed as denoting anything. (Merrill, 2009, p. 14, punctuation added.)

From this passage, and from the absence in Merrill's writings of anything to the contrary, we infer that Merrill, too, is an ostrich nominalist.

Consider now this sentence from his "Ontological Realism" (Merrill, 2010, p. 92):

The point is that while for the metaphysical realist (of the Smith–Ceusters school) a fundamental task of the scientist must be to ask what universals exist, for the anti-realist this is replaced by the much more sensible (and obviously empirical) task of determining what predicates ('loves pizza', 'has the flu', 'is a smoker', 'is a non-smoker', etc.) should be introduced into our scientific language in order to formulate our theories and test them in the empirical world – and which of those predicates we should retain in our language as a consequence of such testing.

How, on the ostrich perspective, could such testing be made intelligible? Let us suppose (somewhat counter-intuitively, given what we know about how scientists work) that a given group of scientists is attempting to determine empirically whether to include the predicate 'has the flu' in their scientific language. How do they do this? By investigating, presumably, whether there are entities in reality which satisfy the predicate 'has the flu'. And how do they do this? Presumably by finding out whether, say, entity Jim satisfies this predicate. And how do they do this? By determining whether the sentence 'Jim has the flu' is true. And how do they do this? By examining whether 'Jim' satisfies the predicate 'has the flu'. And how do they do this? By determining whether the sentence 'Jim has the flu' is true. And so on, *ad indefinitum*.

Merrill seeks to climb out of this circle in the following passage:

we know what it means to have the flu. We can describe tests for determining such a diagnosis and describe clear clinical (empirical) conditions pertaining to the flu and those who have it. The flu universal does not make an appearance. (Merrill, 2010, p. 92.)

But how could it be that we can determine that something, the flu, is *had*, now by this patient, now by that patient, if there are no repeatable somethings (however the latter are to be understood from the metaphysical point of view)? How, if there are no repeatable somethings, could there be *tests* which can be described in a uniform way and reliably applied, now to this patient, now to that patient, to determine whether either has something that would in both cases be referred to as 'the flu'? And how could there be *diagnoses* and *conditions* which share in common that they pertain to the *flu*? How, more generally, is Merrill to do justice to the use of general terms as the *subjects* of true sentences formulating scientific discoveries as, for example in:

(N) The H1N1 virus causes influenza?

As Summerford (2003) argues, 'If the nominalist is going to reject universals, then he must demonstrate that the use of these terms does not involve countenancing such entities', and nominalists have thus far failed to provide a satisfactory demonstration of how this is to be achieved. The one approach which still attracts significant numbers of adherents views general terms such as those appearing in (N) as referring to sets or collections, in effect by identifying universals with their extensions, which is to say with the set or collection of their instances. How, then, to address the problem turning on the fact that multiple putatively very different universals might conceivably have identical extensions in this the actual world? The favored answer to this question (deriving, in its most influential version, from Lewis, 1986) is to view the extensions in question as including members not merely among the

actual, but also among the merely possible, physical individuals – such as, say, Nicola Guarino’s thousandth child. This however creates further problems, not merely because it makes the favored set-theoretic referents for general terms appear (as some might say) *curiously and curiously* the more closely they are scrutinized, but also because it threatens to make the treatment of such terms embarrassingly remote from the scientific and computational needs of, say, biologists.

5.4. First-Order Logic with Universal Terms (FOLWUT)

There is a further problem for the predicate-based approach when received FOL is used, namely that it might leave us in a position where certain needed logical inferences will not be able to be drawn. Consider, to illustrate the point, an assertion concerning some portion a of cell protein extract. From

(O) a was incubated for 10 min.

We can infer:

(P) a was incubated.

One way of treating (O) in received FOL would yield:

(O*) was_incubated_for_10_minutes (a),

a logically not further analyzable sentence, again of the form ‘ $F(a)$ ’, where ‘ F ’ is the predicate and ‘ a ’ is a constant term referring to an individual object to which the predicate ‘ F ’ is applied. From (O*), however (and non-logicians among our readers will be shocked by this), we cannot infer logically the regimented counterpart of (P), namely:

(P*) was_incubated(a).

Famously, this is the problem of the logical analysis of sentences involving adverbial modifications. This problem is commonly seen as having been solved by Ramsey (1978) and Davidson (1980), who recognized that sentences such as (O) are properly to be treated as equivalent to sentences involving existential quantification over events, along the lines of:

(P**) $(\exists e)((\text{instantiates}(e, \text{incubation_event}) \ \& \ \text{participates}(a, e) \ \& \ \text{duration}(e, 10 \text{ minutes}))$

or in other words there is some incubation event e in which a participates and which is of duration 10 minutes.

The inference to (O), which is now regimented as having the form:

(O**) $(\exists e)((\text{instantiates}(e, \text{incubation_event}) \ \& \ \text{participates}(a, e))$

is then a simple matter of conjunction elimination. In this way, we note, needed inferences are secured by appeal to the *representational element* central to the realist methodology.

The importance of Ramsey and Davidson’s work is that they opened the way to relaxing the restrictions of received FOL by allowing terms (the ‘ a ’ in ‘ $F(a)$ ’) to refer not merely to individual *objects* but also to *events*. The realist methodology for ontology development outlined in the second half of Smith (2005) generalizes the Ramsey–Davidson theory by taking this relaxation still further, by allowing terms in FOL to refer not merely to events – a generalization allowed also by SUMO – but also to entities in other categories distinguished in the BFO ontology, above all to entities in the category of *dependent continuant*.

Consider, for example, the sentence:

(Q) Werner has a headache.

On the account of predication dictated by received FOL, such a sentence is to be understood in the same way as ‘Teco is a bonobo’. It has the ‘F(a)’ form:

(R) has_a_headache(Werner).

(R) is true, again, if and only if the subject (Werner) satisfies the predicate (has_a_headache), whereby (R) clearly respects the received FOL rule that it contains only terms referring to individual objects. From the clinico-ontological perspective, however, this rule will pose problems, for it means that (Q) does not allow the inference to, for example,

(S) there is a headache which Werner currently has,

or in symbols:

(Q*) $(\exists x)((\text{instantiates}(x, \text{headache}) \ \& \ \text{inheres}(x, \text{Werner}))$.

Moreover, received FOL will not allow us to assert, for example, that the headache referred to in (S) *has lasted for two hours*, or *is being treated by taking aspirin gum*. Indeed, received FOL allows no reference to disease-entities – to your influenza, or my sinusitis – of any kind; rather, it requires us always to reformulate our statements about such entities as statements about individual objects such as the organisms which are their bearers.

The version (dialect) of FOL that we propose – called ‘FOLWUT’, for: ‘first-order logic with universal terms’ – is designed to resolve such matters by allowing the ways clinicians and others refer to entities such as diseases to be captured using terms in FOL along lines illustrated already in (Q*). But it goes further in allowing terms in FOL to refer not only to independent and dependent continuant particulars and to occurrent particulars, but also to *universals* in all of these categories (Smith, 2005).

FOLWUT thereby departs from received FOL in two ways. First, it expands the repertoire of types of entities to which the *terms* of FOL can refer. At the same time, it radically restricts the family of allowable *predicates*, eliminating all predicates of the usual sort (‘is a man’, ‘is an HMG-CoA reductase inhibitor’ and so forth), and admitting instead only a small number of formal predicates, including two-place (relational) predicates of the sorts described in the Relation Ontology – all of them predicates which, like the formal tie of identity ‘=’, come with fixed interpretations.

Such relational predicates will include, on the level of instances, suitably temporally indexed versions of:

Part_of(x, y), for: individual x is part of individual y

Member_of(x, y), for individual x is a member of individual collection y

Inheres(x, y), for: individual x inheres in individual y

Precedes(x, y), for: individual process x precedes individual process y

Has_Participant(x, y), for: individual thing y participates in individual occurrent x

Has_Agent(x, y), for: individual thing y is agent of individual occurrent x

Realizes(x, y), for: individual process x realizes individual function y .

On the level of universals or types it will include, for example:

is_a(x, y), for: every instance of universal x is part of some instance of universal y

part_of(x, y), for: every instance of universal x is part of some instance of universal y

and finally, bridging the two levels of particulars and universals, it will include relations such as:

Inst(x, y), for: individual x instantiates universal y

Extension(x, y), for: individual collection x is the extension of the universal y

and so on.

The consequence of generalizing the scope of allowed referents for terms in FOL to include also universals is that it brings the possibility of simulating, within an entirely traditional FOL framework, some of the expressive possibilities of second order logic. In particular, we can define, in terms of the instance-instance relations listed above, type-level relations such as *is_a* and *part_of* in ways that are useful not only for ontology-based reasoning but also for ensuring that the relations in question are used by those engaged in the construction of ontologies in ways which avoid certain hitherto common errors (Bittner & Donnelly, 2007; Donnelly et al., 2006). And then, exactly as Merrill would require, the result is a framework in which *predicates do not represent*, and which is governed by standard predicate-logical semantics.

5.5. General terms in scientific hypotheses

We recall that the *principle of instantiation* is formulated only for the case of reference ontologies (and thus of ontologies created in support of settled science). Matters ontological will be more complicated in areas of non-settled science, where there may be multiple camps of experts, and where the appropriate ontological analysis of the very experiments used to test given hypotheses may be subject to dispute. Ontologies may then provide a supporting role in the testing of the relevant hypotheses; however, it is not up to the authors of reference ontologies to pick sides in such disputes; rather this is a decision that should wait for science.

Further issues are raised by our acceptance of the *principle of instantiation*. This principle is designed to ensure that the users of the realist methodology see types or universals not as entities in some special realm that is beyond the reach of empirical observation, but rather squarely within the world of what happens and is the case, entities with which experimental data is associated.

Sometimes, of course, general terms are used by scientists to designate entities (or purported entities) postulated in areas where science is not yet settled, as for example in the case of the Higgs boson (Dumontier & Hoehndorf, 2010). Merrill is right to insist (with Smith, 2006a) that there is a role for ontologies ‘to aid in formulating the hypotheses that later become laws within theories’. Here, clearly, the *principle of instantiation* does not apply. As concerns the other elements of the realist methodology, however – and contrary to what Merrill (2010) and Dumontier and Hoehndorf (2010) argue – ontologies following realist principles are still able to be developed to fulfill this role. The information artifacts in question will not, at least initially, be incorporable into reference ontologies recommended for general use. But they may have a significant practical role to play nonetheless in helping the relevant scientific hypotheses to become part of established science.

According to Merrill, a realist who is faced with a Higgs type of case would either (1) need to wait before beginning the process of ontology building in the relevant areas until the needed universals had ‘emerged’ or (2) require ‘a theory of meaning’ – and thus a (non-realist) theory of ontology – ‘that does not require the Referentialist Assumption’. Case (1) would cripple the realist methodology from a practical point of view. In case (2), realism itself would be sacrificed for at least some portions of ontology building, thereby potentially re-opening the problems flowing from older, concept-based approaches to ontology development.³²

In fact, however, our solution is much more straightforward, and rests on the recognition that language can clearly still be used to communicate – in some sense – even where putative referring expressions fail in their reference. Some people assert their beliefs in the existence of unicorns. All such beliefs are false. But the beliefs *exist* just as do other beliefs; they can be communicated; and they can also be represented, as what we have called ‘level 2 entities’ (Smith et al., 2006), in realist ontologies, created, perhaps, for purposes of supporting psychiatric research.

The case of psychiatry reminds us that the issues raised by non-referring terms apply just as much to singular terms as they do at the level of the general terms appropriate to ontologies. Let us suppose, for example, that a psychiatric patient begins to express beliefs in something he calls ‘Murther’. For the moment, we do not know whether ‘Murther’ refers to some entity or whether it is, like ‘unicorn’, merely the expression of some fantasy. Until this matter is settled the psychiatrist, in compiling his clinical record for the patient, can avail himself of the facility incorporated in the Referent Tracking paradigm, whereby instance unique identifiers can be *reserved* for candidate particulars whose existence is not yet settled, for example, when an order to obtain X-ray studies on some patient has been entered into the hospital order system today, and identifiers are needed already in advance of the radiographs that will exist only tomorrow. Such identifiers will lose their ‘reserved’ status once the entities in question have been confirmed to exist; and they will be immediately declared obsolete should it ever be confirmed that the putative entities in question do not and will never exist. The formal mechanisms are introduced in Ceusters (2007). We have recommended that analogous mechanisms be formulated for application ontologies, incorporating also new *evidence codes* to indicate that assertions containing the terms in question are, for different reasons, problematic.³³

By employing such mechanisms, application ontologies following realist principles can be developed even where general terms are being used before the existence of corresponding instances has been confirmed. The terms in question would need only to be provided with provisional identifiers for purposes of ontological reasoning support.

The proposal thus conforms well with the strategy already implemented in the chemistry domain, an area which Dumontier and Hoehndorf (2010) argue might somehow not be well served by the realist approach. Consider, for example the model followed by IUPAC, the International Union of Pure and Applied Chemistry, in its treatment of elements. There, a formal name is given to an element only after evidence has been presented that it has been created in the lab and this evidence has been verified through a rigorous process. In the meantime, IUPAC creates provisional names for those elements hypothesized to exist, but the latter are not included in the pertinent reference ontology (i.e., the Periodic Table) until officially proven. At the same time, of course, pharmaceutical and other organizations are developing the equivalent of application ontologies to support their planning processes in which terms may be reserved, for example, for chemical substances that have not yet been synthesized.

The problems identified in the above do not pertain specifically to ontologies and to the role of the general terms therein. They pertain quite generally to uses of terms to (putatively)

³²Some might suppose that there is a case (3), involving a hybrid approach that uses *universals* for scientifically established entities and *meanings* or *concepts* for Higgs-type cases. One problem with such an approach, however, is that it leaves the ontologist with an incoherent account of the referents of such terms during the transition from speculative to settled usage, as for example in the case of new diseases at the stage when patients are already affected but the diseases themselves have not yet been incorporated into settled diagnostic science.

³³http://sourceforge.net/mailarchive/message.php?msg_name=20100719152134.E60EE207A9%40mweb2.acsu.buffalo.edu, last accessed August 10, 2010.

refer to what does not exist (Kroon, 1992). In the realm of particulars such terms are often used in the context of planning – for example, in the naming of babies not yet even conceived. And for the formal regimentation of processes of this sort it appears that the appeal to some sort of possible worlds approach à la Lewis is what is required.

We have argued that a reference ontology is analogous to a settled scientific theory (Smith, 2008). Developing such an ontology presupposes an intention on the part of the developer to represent some configuration of repeatable structures in reality in a way that conforms to the current content of the relevant parts of science. There are of course many ontology-like artifacts which rest on different goals. Some might develop application ontologies to capture, for example, the content of Klingon science. Some might develop application ontologies in the service of the history of science to represent entities postulated by Earth-based scientific theories that have long since been falsified. But such artifacts, we believe, must be sequestered from the reference ontologies recommended for general use in support of scientific research.

Merrill argues that a further problem arises for our views in the case of those general terms whose referents have not yet been confirmed, because: ‘If hypotheses containing such terms can be regarded as meaningful (and they must if they are to be tested), then it cannot be required that the terms in them denote universals’ (Merrill, 2010, p. 87). Here, too, we believe, a response assigning the relevant term to an appropriate application ontology will be quite sufficient.

6. BFO, DOLCE, SUMO, Cyc

Of the four leading upper-level ontologies in the public domain – BFO, DOLCE, SUMO and OpenCyc – BFO is in one respect more closely tailored to the needs of scientist users. This is because it is a strict upper ontology, which means that it does not contain its own representations of physical, chemical, biological, psychological, social or other types of entities which would properly fall within the domains of the special sciences. This reflects the fact that BFO was developed as a very small representational artifact with the narrowly focused task of providing an upper ontology that could be used to support the integration of multiple heterogeneous ontologies developed for purposes of scientific research.

DOLCE (Gangemi et al., 2002) is from the point of view of numbers of users a very successful upper-level ontology, and it has been applied in a number of projects in biomedical^{34,35} and other scientific domains. DOLCE and BFO in fact grew out of a common philosophical orientation, and thus BFO overlaps with parts of DOLCE’s top level and is in close conformity with the DOLCE-associated OntoClean methodology (Guarino and Welty, 2002). But DOLCE has chosen a strategy, different from that of BFO, focusing on what it calls ‘linguistic and cognitive engineering’. This means that its coverage domain includes the putative objects of mythology (leprechauns, for example) or fiction (instances of pneumonia in 19th century Russian novels) and thus that, unlike BFO, it relies on an ontology of possible worlds. We do not believe that this makes DOLCE stronger from the perspective of providing support for the development of reference ontologies to serve the needs of scientific researchers.

SUMO, too, has proved to have considerable value as an upper-level ontology for certain purposes (Niles & Pease, 2001).³⁶ Unfortunately the fact that it contains its own tiny biology (‘protein’, ‘crustacean’, ‘body-covering’, ‘fruit-Or-vegetable’) means that it cannot

³⁴<http://neuroscientific.net/index.php?id=43>, last accessed June 30, 2010.

³⁵<http://www.imbi.uni-freiburg.de/aneurist/ontology/>, last accessed June 30, 2010.

³⁶It also incorporates certain elements contributed by Smith: <http://suo.ieee.org/SUO/Ontology-refs.html>, last accessed June 30, 2010.

support the strategy of downward population that has proved so useful to scientists in the case of BFO, since biologists are unlikely to find SUMO's definitions (and selection) of biological terms acceptable, and they will find problematic the absence in SUMO of anything like the BFO category of dependent continuant (for particulars such as Werner's headache, Mary's hypertension or Bruno's osteoarthritis (Scheuermann et al., 2009)).

Merrill has a number of positive things to say about Cyc in his "Ontological Realism". Both DOLCE and SUMO seem to us, however, to be much more coherent as ontologies for scientific purposes than the upper level of Cyc, which is marred not least by the fact that it is associated at lower levels with very many terms and definitions which, because of Cyc's primary focus on formalizing what it calls 'common sense knowledge', deviate significantly from the terms and definitions favored by scientists. (The children of Cyc's *partially tangible thing*, for example, include both *diisopropyl methylphosphonate* and *pay e-mail provider*.) From our present perspective, however, Cyc's primary problem turns on the fact that (like the UMLS) it does not strive for consistency among the various 'microtheories' which form its parts. Hence the very goal of creating a single consistent suite of interoperable ontologies which would capture the terminological content of biomedical science – which is from our point of view the only coherent strategy for achieving ontology-mediated data integration in the domain of the life sciences – is undermined by Cyc's own paraconsistent logical structure.

DOLCE and SUMO are of signal importance for our argument here, however, because, like BFO, both are constructed around (overwhelmingly) single inheritance taxonomies (*is_a* hierarchies) consisting of singular nouns representing what in BFO and DOLCE (Masolo et al., 2002) are called 'universals' and in SUMO 'classes'.³⁷ In each case, the generic entities which form the focus of the ontologies are said to have instances in the realm of particulars. In each case the generic entities are governed by the sparse theory of universals outlined in our discussion above.

Whether all of this applies to the Cyc knowledge base also is, alas, not easy to ascertain from its documentation. But it is in any case clear – and surely significant – that at least three of the four leading upper-level ontologies rest on views concerning the relation between general terms and universals of just the sort that Merrill finds so objectionable.

Another design choice shared in common by BFO, DOLCE and SUMO is the acceptance of a dichotomy between *continuants* and *occurrents*. Philosophers have argued back and forth for some two thousand years over the question whether this dichotomy is truly such as to represent the fundamental architecture of the reality that (as we would now say) is described by science. Such arguments continue to be pursued at length by distinguished figures in the ontology field such as John Sowa, who sees the continued existence of philosophical communities with opposing views on this and similar matters as justification for his own long-standing campaign against the very project of a consistent, formalized, upper-level ontology of the sort whose widespread adoption is, in our eyes, the *sine qua non* of effective ontology coordination in a large multidisciplinary area such as biomedicine. Rather, Sowa favors a Cyc-like approach to ontology, in which the sparse theory of universals is abandoned in favor of the acceptance of unrestricted Boolean combinations – an approach which Sowa himself describes as

³⁷Classes are elucidated in the SUMO documentation as follows: 'Classes differ from Sets in two important respects. First, Classes are not assumed to be extensional. That is, distinct Classes might well have exactly the same instances. Second, Classes typically have an associated "condition" that determines the instances of the Class. So, for example, the condition "human" determines the Class of Humans'. See <http://www.ontologyportal.org/translations/SUMO.owl.txt>, last accessed June 30, 2010.

a way of formalizing a structure of microtheories to accommodate an open-ended, possibly inconsistent *knowledge soup*, [organized] in an infinite lattice, which would be rich enough to include any possible language game that any finite reasoner (human, computer, or extraterrestrial) could ever invent. (Sowa, 2006, emphasis added.)

Sowa's position as here formulated may be appropriate as concerns ontologies designed for purposes other than those of concern to us here. Viewed against the background of the issues discussed in the present paper, however, his position appears to us to be analogous to that of one who would hold that, because there are equal and opposite arguments in favor of both *driving on the left* and *driving on the right*, we should avoid the imposition of any 'rigid formalized traffic system' and foster, instead, an open-ended, possibly inconsistent *traffic soup*, which would be rich enough to include any possible traffic game that any finite driver (human, computer or extraterrestrial) could ever invent.

For our part, we embrace instead a view based on the idea that, for coordination purposes, path dependence may be a *good thing*. If there are 2, or 7, different and equally effective ways to reach a certain end, and if we know that significant benefits will accrue from choosing just one of these ways, then we should choose one way and vigorously propagate this choice. In ontology, specifically, this means that we accept the

Ontological traffic law principle: Ontological standards, including a common upper-level ontology and standards governing syntactical uniformity, are indispensable to every successful large-scale ontology development initiative, and this is so even if they are selected arbitrarily provided they enjoy widespread assent among those working in the relevant research community.

One example of such a traffic law, which has been executed with some success and, we believe, some measurable benefit by the GO and its sister ontologies within the OBO Foundry (Smith et al., 2007), is the law according to which all terms within an ontology should be nouns and noun-phrases that are *singular in number*. (This purely syntactic law is in fact inspired by our view according to which ontologies should be viewed as consisting of representations of types or universals, but its implementation need clearly involve no reference to this view.) Another example is the law which asserts that all terms in an ontology should be traceable via *is_a* relations to the relevant ontology root node. Further examples of such laws have been codified by the OBO Foundry in the form of an evolving set of principles for ontology development in the biological and medical domains, some of them focusing on governance. The first ten of these principles, first promulgated in April 2006,³⁸ have proved to be of value to ontology developers seeking guidance on how most effectively to create ontologies in such a way as to maximize consistency with other OBO ontologies. Further principles are currently under review by the OBO Foundry with a view to their adoption in the future.

These principles are interesting, since some of them are treated by Merrill as figures of fun, and some of them as dangers to the advance of science. Under the first heading, Merrill sees some of the principles of ontological realism roughly along the lines of 'How could such a strange amalgam of Aristotelico-Australian philosophical ideas possibly have import for the workings of serious scientific research?' In fact, however, since the authors of this communication first began to collaborate in 2002, our ontology development methodology has been driven by needs and concerns not of philosophers, but rather of scientists building systems in areas such as hospital adverse event reporting (Ceusters et al., 2009, 2009a, 2009b), salivaomics (Ai et al., 2010), or the diagnosis and treatment of Methicillin-resistant *Staphylococcus aureus* (Goldfain et al., 2010).

³⁸<http://www.obofoundry.org/crit.shtml>, last accessed June 30, 2010.

Under the second heading (dangers to the advance of science), Merrill is concerned that the OBO Foundry principle of modularity – according to which there should be one ontology for each domain that is recommended for general use in realizing the purposes of the Foundry – might harbor a view according to which for every scientific domain there is or will be exactly one true theory, a view which could have detrimental consequences in constraining the flexibility that is indispensable to scientific advance. As we have argued at length (Smith et al., 2007), however, the OBO Foundry is not attempting to restrict the ontologies people can build. Rather, it is attempting, as an experiment, to create a suite of ontology artifacts built around a small set of high quality, interoperable, non-overlapping reference ontologies following certain principles. All of those involved in the Foundry initiative recognize that it is vital to the success of the Foundry that it is always open to, and can only benefit from, both (1) criticism from the outside – on the basis of the assumption that no Foundry resource will ever exist in a form that cannot be further improved, and (2) competitor initiatives, both at the level of single ontologies and at the level of the Foundry as a whole.

7. Conclusion

At one point Merrill (2010, p. 105) asserts that our approach ‘*is neither science nor philosophy*’, and in this he hits the nail exactly on the head. For in propagating the realist methodology we are indeed engaging in a novel interdisciplinary activity that involves elements of both of these, and also of computer science, politics, community organizing, sociology, logic, and other black arts. Merrill himself, however, draws a slightly different conclusion. For him, our approach – because it involves reference to those damned “universals” – is ‘ideology’ through and through, and ‘hence in the final analysis unscientific’. Let us grant him, in the interests of eirenic compromise, that there *is* an element of ideology involved in our work. Coordinated ontology development across a large scale is so difficult that we are happy to draw on any means that will help us to achieve our ends. But then at the same time we submit that there is an equal and opposite admixture of ideology on Merrill’s side also – an ideology deriving from the School of Nominalism.

For this reason too, therefore, we would welcome a systematic effort on Merrill’s part to create and disseminate a strategy for ontology development that can be *certified to be general term free*. If such a strategy were to gain traction amongst biologists, to the point where Merrill himself were able to point to evidence of clear practical advantages over the realist approach, then we would of course switch our adherence immediately. Strangely, though, we cannot shake off our conviction that Merrill himself, were he to find himself in an analogous situation,³⁹ would *not* switch over to our side.

Acknowledgments

We are grateful to Gary Merrill for giving us this opportunity to clarify our views. We thank also Colin Batchelor, Randall Dipert, Albert Goldfain, Janna Hastings, William Hogan, Ingvar Johansson, Michael McGlone, Chris Mungall, Peter Robinson, Stefan Schulz, David Osumi-Sutherland, Alan Ruttenberg, Frederic Tremblay, Neil Williams and the participants in the obo-discuss email discussion at <http://tinyurl.com/34lacvy>, for valuable suggestions. The work on this paper was partially supported by the National Institutes of Health through the NIH Roadmap for Medical Research, Grant 1 U 54 HG004028 (National Center for Biomedical Ontology) and also by Grant R21LM009824 from the National Library of Medicine. The content of this paper is solely the responsibility of the authors and does not necessarily represent the official views of the National Library of Medicine or the National Institutes of Health.

³⁹It may be that Merrill *is already in this analogous situation*, for instance given the statistics assembled in Bodenreider (2008), which documents a significant fall-off in citations of the UMLS by clinical and biological researchers in recent years and a countervailing rise in usage of the GO (measured as a percentage of all PubMed citations pertaining to ontology).

References

- Abazov VM, et al. Search for Higgs boson production in dilepton and missing energy final states with 5.4 fb^{-1} of $p\bar{p}$ collisions at 2 1.96 TeV. *Physical Review Letters*. 2010; 104(061804)
- Ai J, Smith B, Wong D. Saliva ontology: An ontology-based framework for a salivaomics knowledge base. *BMC Bioinformatics*. 2010; 11:302. [PubMed: 20525291]
- Arighi C, Liu H, Natale D, Barker W, Drabkin H, Hu Z, Blake J, Smith B, Wu C. TGF-beta signaling proteins and the protein ontology. *BMC Bioinformatics*. 2009; 10(Suppl. 5) Art. No. S3.
- Armstrong, DM. Universals and Scientific Realism. Nominalism and Realism (Vol. 1). *A Theory of Universals* (Vol. 2). Cambridge University Press; Cambridge: 1978.
- Armstrong DM. Against ostrich nominalism: a reply to Michael Devitt. *Pacific Philosophical Quarterly*. 1980; 61:441.
- Armstrong DM. In defence of structural universals. *Australasian Journal of Philosophy*. 1986; 64(1): 85–88.
- Armstrong, DM. *Universals: An Opinionated Introduction*. Westview Press; Boulder, CO: 1989.
- Armstrong, DM. Universals as attributes. In: Loux, MJ., editor. *Metaphysics: Contemporary Readings*. 2nd edn.. Routledge; New York: 2008.
- Batchelor, C.; Bittner, T.; Eilbeck, K.; Mungall, C.; Richardson, J.; Knight, R.; Stombaugh, J.; Zirbel, CL.; Westhof, E.; Leontis, NB. The RNA Ontology (RNAO): an ontology for integrating RNA sequence and structure data; Proceedings of the International Conference on Biomedical Ontologies; Buffalo, NY: University at Buffalo. 2009; p. 7-10.
- Bittner T, Donnelly M. Logical properties of foundational relations in bio-ontologies. *Artificial Intelligence in Medicine*. 2007; 39:197–216. [PubMed: 17428644]
- Bittner, T.; Smith, B. Vagueness and granular partitions. In: Welty, C.; Smith, B., editors. *Formal Ontology and Information Systems*. ACM Press; New York: 2001. p. 309-321.
- Bodenreider, O. Yearbook of Medical Informatics. Schattauer; Stuttgart: 2008. Biomedical ontologies in action: role in knowledge management, data integration and decision support; p. 67-79.
- Bodenreider, O.; Smith, B.; Burgun, A. The ontology-epistemology divide: a case study in medical terminology. In: Varzi, A.; Vieu, L., editors. *Formal Ontology and Information Systems; Proceedings of the Third International Conference (FOIS 2004)*; Amsterdam: IOS Press. 2004; p. 185-195.
- Bourget, D.; Chalmers, D. The PhilPapers surveys: results, analysis and discussion. 2009. Available at: <http://philpapers.org/surveys/>
- Brown, DE. *Human Universals*. McGraw-Hill; New York: 1991.
- Buszkowski, W.; Marciszewski, W.; Benthem, JV., editors. *Categorial Grammar*. John Benjamins; Amsterdam: 1988.
- Carnap, R. *Der logische Aufbau der Welt*. Felix Meiner; Leipzig: 1928. English translation by R.A. George, *The Logical Structure of the World. Pseudoproblems in Philosophy*. University of California Press, 1967
- Carnap R. Empiricism, semantics, and ontology. *Revue Internationale de Philosophie*. 1950; 4:20–40.
- Cavalli-Sforza LL. Genes, peoples, and languages. *Proceedings of the National Academy of Sciences*. 1997; 94:7719–7724.
- Ceusters, W. Dealing with mistakes in a referent tracking system. In: Hornsby, KS., editor. *Proceedings of Ontology for the Intelligence Community (OIC)*. Columbia, MA: November 28–29. 2007 p. 5-8.
- Ceusters W. Applying evolutionary terminology auditing to the Gene Ontology. *Journal of Biomedical Informatics*. 2009; 42(3):518–529. [PubMed: 19162233]
- Ceusters, W.; Capolupo, M.; Devlies, J. D4.2 – RAPS Domain Ontology (M12 Version). Background materials and methodology used to develop the domain ontology for risks against patient safety. 2009a. Available at: <http://www.referent-tracking.com/RTU/sendfile/?file=ReMINE-D4-2.pdf>
- Ceusters, W.; Capolupo, M.; Devlies, J. D4.3 – RAPS Application ontology (Version 1). Background materials and methodology used to develop application ontologies for risks against patient safety. 2009b. Available at: <http://www.referent-tracking.com/RTU/sendfile/?file=ReMINE-D4-3.pdf>

- Ceusters W, Capolupo M, Smith B, De Moor G. An evolutionary approach to the representation of adverse events, *Medical Informatics Europe (MIE 2009)*, Sarajevo. Studies in Health Technology and Informatics. 2009; 150:537–541. [PubMed: 19745369]
- Ceusters, W.; Smith, B. A realism-based approach to the evolution of biomedical ontologies; Proceedings of the Annual AMIA Symposium; Washington, DC. 2006a; p. 121-125.
- Ceusters W, Smith B. Strategies for referent tracking in electronic health records. *Journal of Biomedical Informatics*. 2006b; 39:362–378. [PubMed: 16198639]
- Ceusters, W.; Smith, B. Proceedings of Medinfo. Cape Town, South Africa: 2010. A unified framework for biomedical terminologies and ontologies; p. 1050-1054.
- Ceusters, W.; Smith, B.; Flanagan, J. Ontology and medical terminology: why Description logics are not enough; Proceedings of the Conference: Towards an Electronic Patient Record (TEPR 2003); Boston, MA: Medical Records Institute. 2003; CD-ROM publication
- Ceusters W, Smith B, Goldberg L. A terminological and ontological analysis of the NCI Thesaurus. *Methods of Information in Medicine*. 2005; 44:498–507. [PubMed: 16342916]
- Ceusters, W.; Spackman, K.; Smith, B. Would SNOMED-CT benefit from realism-based ontology evolution?; Proceedings of the Annual Symposium of the American Medical Informatics Association; Chicago, IL. 2007; p. 105-109.
- Ceusters W, Elkin P, Smith B. Negative findings in electronic health records and biomedical ontologies: a realist approach. *International Journal of Medical Informatics*. 2007; 76:326–333.
- Cocchiarella NB. Conceptual realism and the nexus of predication. *Metalogicon*. 2003; 16(2):45–70.
- Comrie, B. *Language Universals and Linguistic Typology: Syntax and Morphology*. 2nd edn.. University of Chicago Press; Chicago, IL: 1989.
- Courtot M, Gibson F, Lister AL, Malone J. MIREOT: the minimum information to reference an external ontology term. *Applied Ontology*. 2011 to appear.
- Davidson, D. *Essays on Actions and Events*. Clarendon; Oxford: 1980.
- Diehl AD, Deckhut A, Blake JA, Cowell LG, Gold ES, Gondré-Lewis TA, Masci AM, Meehan TM, Morel PA, Nijnik A, Peters B, Pulendran B, Scheuermann RH, Yao QA, Zand MS, Mungall CJ. Hematopoietic cell types: prototype for a revised Cell Ontology. *Journal of Biomedical Ontology*. 2010 to appear.
- Donnelly M, Bittner T, Rosse C. A formal theory for spatial reasoning in biomedical ontologies. *Artificial Intelligence in Medicine*. 2006; 36:1–27. [PubMed: 16249077]
- Dumontier, M.; Hoehndorf, R. Realism for scientific ontologies; Proceedings of the 6th International Conference on Formal Ontology in Information Systems (FOIS); Toronto. 2010; p. 387-399.
- Gangemi, A.; Guarino, N.; Masolo, C.; Oltramari, A.; Schneider, L. *Knowledge Engineering and Knowledge Management: Ontologies and the Semantic Web*. LNCS. Vol. 2473. 2002. Sweetening ontologies with DOLCE; p. 223-233.
- Gene Ontology Consortium. Gene Ontology: tool for the unification of biology. *Nature Genetics*. 2000; 25(1):25–29. [PubMed: 10802651]
- Ghiselin, MT. *Metaphysics and the Origin of Species*. State University of New York Press; Albany: 1997.
- Givón T. Bumping into Joe, repeatedly: Joseph Greenberg the theorist. *Linguistic Typology*. 2002; 6(1):8–16.
- Goldfain A, Smith B, Cowell LG. Towards an ontological representation of resistance: The case of MRSa. *Journal of Biomedical Informatics*. 2010 to appear.
- Goodman, N. *The Structure of Appearance*. Harvard University Press; 1951.
- Greenberg, JH., editor. *Universals of Languages*. MIT Press; Cambridge, MA: 1963.
- Grenon P, Smith B. SNAP and SPAN: towards dynamic spatial ontology. *Spatial Cognition and Computation*. 2004; 4(1):69–103.
- Gruber T. Toward principles for the design of ontologies used for knowledge sharing. *International Journal Human-Computer Studies*. 1992; 43:907–928.
- Guarino N, Welty C. Evaluating ontological decisions with OntoClean. *Communications of the ACM Archive*. 2002; 45(2):61–65.

- Haendel, M.; Neuhaus, F.; Sutherland, D.; Mejino, JLE., Jr.; Mungall, C.; Smith, B. CARO: The Common Anatomy Reference Ontology. In: Burger, A.; Davidson, D.; Baldock, R., editors. *Anatomy Ontologies for Bioinformatics: Principles and Practice*. Springer; New York: 2008. p. 327-349.
- Hill DP, Blake JA, Richardson JE, Ringwald M. Extension and integration of the Gene Ontology (GO): combining GO vocabularies with external vocabularies. *Genome Research*. 2002; 12(12): 1982–1991. [PubMed: 12466303]
- Hill DP, Smith B, McAndrews-Hill MS, Blake JA. Gene Ontology annotations: what they mean and where they come from. *BMC Bioinformatics*. 2008; 9(Suppl. 5):S2. [PubMed: 18460184]
- Hogan, WR. What's in an 'is a' link?; Proceedings of the First International Conference on Biomedical Ontology; Buffalo. 2009; p. 170
- Hogan WR. Why the Unified Medical Language System is not an ontology, MS. 2010
- Hogan WR. Towards an ontological theory of substance intolerance and hypersensitivity. *Journal of Biomedical Informatics*. 2010 to appear.
- Holenstein, E. *Roman Jakobson's Approach to Language*. Indiana University Press; Bloomington, IN: 1976.
- Husserl, E. *Logische Untersuchungen*. 2nd edn.. Niemeyer; Halle: 1913/21, 1970. English translation as *Logical Investigations*, by J.N. Findlay. London: Routledge and Kegan Paul
- ISO. *Terminology-Vocabulary (ISO 1087: 1990)*. International Standards Organization; Geneva: 1990.
- ISO. *Text for FDIS 704. Terminology work: principles and methods (ISO/IEC JTC1 SC36 N0579: 1999)*. International Standards Organization; Geneva: 1999.
- Johansson, I. Pattern as an ontological category. In: Guarino, N., editor. *Formal Ontology in Information Systems*. IOS Press; Amsterdam: 1998. p. 86-94.
- Kroon FW. Was Meinong only pretending? *Philosophy and Phenomenological Research*. 1992; 52(3): 499–527.
- Kuhn, TS. *The Structure of Scientific Revolutions*. The University of Chicago Press; Chicago, IL: 1970.
- Kuroda S-Y. A second look at Marty, Husserl, and Chomsky: the significance of the revolution in linguistics. *Tohoku Daigaku Kenkyu Nenpo*. 1997; 47:1–37.
- Lenat D. CYC: a large-scale investment in knowledge infrastructure. *Communications of the ACM Archive*. 1995; 38(11):33–38.
- Lewis, D. *On the Plurality of Worlds*. Blackwell; Oxford: 1986.
- McCarthy, J.; Hayes, PJ. Some philosophical problems from the standpoint of artificial intelligence. In: Meltzer, B.; Michie, D., editors. *Machine Intelligence*. Vol. 4. Edinburgh University Press; Edinburgh: 1969. p. 463-502.
- Masci AM, Arighi CN, Diehl AD, Lieberman AE, Mungall C, Scheuermann RH, Smith B, Cowell LG. An improved ontological representation of dendritic cells as a paradigm for all cell types. *BMC Bioinformatics*. 2009; 10:70. [PubMed: 19243617]
- Masolo, C.; Borgo, S.; Gangemi, A.; Guarino, N.; Oltramari, A. WonderWeb Deliverable D18: Ontology Library (Final). 2002. Available at: <http://wonderweb.semanticweb.org/deliverables/documents/D18.pdf>
- de Matos P, Alcántara R, Dekker A, Ennis M, Hastings J, Haug K, Spiteri I, Turner S, Steinbeck C. Chemical entities of biological interest: an update. *Nucleic Acids Research*. 2010; 38:D249–D254. [PubMed: 19854951]
- Merrill, GH. Engineering a development platform for ontology-enhanced knowledge applications. In: Sharman, R., et al., editors. *Ontologies. A Handbook of Principles, Concepts and Applications in Information Systems*. Springer; New York: 2007. p. 777-822.
- Merrill GH. Concepts and synonymy in the UMLS Metathesaurus. *Journal of Biomedical Discovery and Collaboration*. 2009; 4(7):1–37. [PubMed: 19126221]
- Merrill GH. Ontological realism: methodology of misdirection? *Applied Ontology*. 2010; 5:79–108.
- Morris C. *On the history of the International Encyclopedia of Unified Science*. Synthese. 1960; 12:517–521.

- Mungall CJ. Obol: Integrating language and meaning in bio-ontologies. *Comparative and Functional Genomics*. 2004; 5:509–520. [PubMed: 18629143]
- Mungall CJ, Batchelor C, Eilbeck K. Evolution of the Sequence Ontology terms and relationships. *Journal of Biomedical Informatics*. 2010 to appear.
- Neuhaus, F.; Grenon, P.; Smith, B. In: Varzi, A.; Vieu, L., editors. *A formal theory of substances, qualities, and universals; Formal Ontology in Information Systems: Proceedings of the Third International Conference (FOIS 2004)*; Amsterdam: IOS Press. 2004; p. 49-59.
- Niles, I.; Pease, A. In: Welty, C.; Smith, B., editors. *Towards a standard upper ontology; Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS)*; Amsterdam: ACM Press. 2001; p. 2-9.
- Noy, NF.; McGuinness, DL. *Ontology development 101: a guide to creating your first ontology*, Technical report. 2001. Available at:
http://protege.stanford.edu/publications/ontology_development/ontology101-noy-mcguinness.html
- Pinker, S. *The Blank Slate*. Viking Press; New York: 2002.
- Ramsey, FP. *Foundations*. Mellor, DH., editor. Routledge; London: 1978.
- Quine, WVO. *From a Logical Point of View*. Harvard University Press; Cambridge: 1953. On what there is.
- Rector, AL. *Proceedings of K-CAP*. ACM Press; New York: 2003. Modularisation of domain ontologies implemented in description logics and related formalisms including OWL; p. 121-128.
- Rector AL, Nowlan WA. The GALEN project. *Computer Methods and Programs in Biomedicine*. 1994; 45(1,2):75–78. [PubMed: 7889770]
- Rosse C, Mejino JLE Jr. A reference ontology for biomedical informatics: the Foundational Model of Anatomy. *Journal of Biomedical Informatics*. 2003; 36:478–500. [PubMed: 14759820]
- Rosse, C.; Mejino, JLE, Jr.. *The Foundational Model of Anatomy Ontology*. In: Burger, A.; Davidson, D.; Baldock, R., editors. *Anatomy Ontologies for Bioinformatics: Principles and Practice*. Springer; London: 2007. p. 59-117.
- Scheuermann, RH.; Ceusters, W.; Smith, B. *Toward an ontological treatment of disease and diagnosis; Proceedings of the 2009 AMIA Summit on Translational Bioinformatics*; Washington, DC: AMIA. 2009; p. 116-120.
- Schulz S, Sunitisrivaraporn B, Baader F, Boeker M. SNOMED reaching its adolescence: ontologists' and logicians' health check. *International Journal of Medical Informatics*. 2009; 78(Suppl. 1):S86–S94. [PubMed: 18789754]
- Slater BH. Internal and external negations. *Mind*. 1979; 38(1):588–591.
- Smith, B. Husserl, language and the ontology of the act. In: Buzzetti, D.; Ferriani, M., editors. *Speculative Grammar, Universal Grammar, and Philosophical Analysis of Language*. John Benjamins; Amsterdam: 1987. p. 205-227.
- Smith, B. *Towards a history of speech act theory*. In: Burkhardt, A., editor. *Speech Acts, Meanings and Intentions. Critical Approaches to the Philosophy of John R. Searle*. de Gruyter; Berlin/New York: 1990. p. 29-61.
- Smith, B. *Beyond concepts: Ontology as reality representation; Proceedings of the Third International Conference on Formal Ontology in Information Systems (FOIS 2004)*; Amsterdam: IOS Press. 2004; p. 73-84.
- Smith, B. In: Reicher, ME.; Marek, JC., editors. *Against fantology; Experience and Analysis: Papers of the 27th International Wittgenstein Symposium*; Vienna: The Austrian Ludwig Wittgenstein Society. 2005; p. 153-170.
- Smith B. *From concepts to clinical reality: an essay on the benchmarking of biomedical terminologies*. *Journal of Biomedical Informatics*. 2006a; 39:288–298. [PubMed: 16293444]
- Smith, B. *Against idiosyncrasy in ontology development*. In: Bennett, B.; Fellbaum, C., editors. *Formal Ontology in Information Systems; Proceedings of the Fourth International Conference*; Amsterdam: IOS Press. 2006b; p. 15-26.
- Smith, B. *Ontology (science)*. In: Eschenbach, C.; Gruninger, M., editors. *Formal Ontology in Information Systems; Proceedings of the Fifth International Conference*; Amsterdam: IOS Press. 2008; p. 21-35.

- Smith, B. HL7 watch. 2010. Available at: <http://hl7-watch.blogspot.com/>
- Smith B, Ashburner M, Rosse C, Bard J, Bug W, Ceusters W, Goldberg L, Eilbeck K, Ireland A, Mungall CJ, The OBI Consortium, Leontis N, Rocca-Serra P, Ruttenberg A, Sansone S-A, Scheuermann RH, Shah N, Whetzel PL, Lewis S. The OBO Foundry: coordinated evolution of ontologies to support biomedical data integration. *Nature Biotechnology*. 2007; 25(11):1251–1255.
- Smith B, Brogaard B. A unified theory of truth and reference. *Logique et Analyse*. 2000; 169:49–93. 170.
- Smith B, Ceusters W. Towards industrial-strength philosophy: how analytical ontology can help medical informatics. *Interdisciplinary Science Reviews*. 2003; 28:106–111.
- Smith B, Ceusters W. HL7 RIM: an incoherent standard. *Studies in Health Technology and Informatics*. 2006; 124:133–138. [PubMed: 17108516]
- Smith B, Ceusters W, Klagges B, Köhler J, Kumar A, Lomax J, Mungall CJ, Neuhaus F, Rector A, Rosse C. Relations in biomedical ontologies. *Genome Biology*. 2005; 6(5):R46. [PubMed: 15892874]
- Smith B, Ceusters W, Temmerman R. Wüsteria. *Studies in Health Technology and Informatics*. 2005; 116:647–652. In *Medical Informatics Europe (MIE 2005)*. Geneva. [PubMed: 16160331]
- Smith, B.; Köhler, J.; Kumar, A. Lecture Notes in Bioinformatics. Vol. 2994. Springer; Berlin: 2004. On the application of formal principles to life science data: a case study in the Gene Ontology; p. 79-94. In *Proceedings of DILS 2004 (Data Integration in the Life Sciences)*
- Smith, B.; Kusnierczyk, W.; Schober, D.; Ceusters, W. Towards a reference terminology for ontology research and development in the biomedical domain. In: Bodenreider, O., editor. *Proceedings of KR-MED*. Baltimore, MD: 2006. p. 57-66. Available at: <http://ceur-ws.org/Vol-222>
- Smith, B.; Rosse, C. The role of foundational relations in the alignment of biomedical ontologies. In: Fieschi, M., et al., editors. *Medinfo 2004*. IOS Press; Amsterdam: 2004. p. 444-448.
- Solbrig, HR.; Chute, CG. Concepts, modeling and confusion; *Proceedings of the First International Conference on Biomedical Ontology*; Buffalo, NY. 2009; p. 121-125.
- Sowa, JF. Signs, processes, and language games: foundations for ontology. 2006. Available at: <http://www.jfsowa.com/pubs/signproc.htm>
- Summerford J. Neither universals nor nominalism. *Kinds and the problem of universals*. *Metaphysica*. 2003; 3:101–126.
- Topalis P, Dialynas E, Mitraka E, Deligianni E, Siden-Kiamos I, Louis C. A set of ontologies to drive tools for the control of vector-borne diseases. *Journal of Biomedical Informatics*. 2010 to appear.
- Woodger, JJ. *The Axiomatic Method in Biology*. Cambridge University Press; Cambridge: 1937.
- Wroe, CJ.; Stevens, RD.; Goble, CA.; Ashburner, M. A methodology to migrate the Gene Ontology to a Description Logic environment using DAML+OIL; 8th Pacific Symposium on Biocomputing (PSB); Lihue, Hawaii. 2003; p. 624-636.
- Zalta, E. *Abstract Objects: An Introduction to Axiomatic Metaphysics*. D. Reidel; Dordrecht: 1983.
- Zhang M, Chen W, Smith SM, Napoli JL. Molecular characterization of a mouse short chain dehydrogenase/reductase active with all-trans-retinol in intact cells, mRDH1. *Journal of Biological Chemistry*. 2001; 276(47):44083–44090. [PubMed: 11562362]




Fig. 1. Proposed Elementary Particle Ontology according to the Standard Model (*Anno* 2009).

relationship	CONTINUANT			OCCURRENT
	INDEPENDENT	DEFINITION	DEFINITION	
ORGAN AND ORGANISM	Organism (NCBI Taxonomy)	Anatomical Entity (MA, CARC)	Organ Function (PATO)	Biological Process (GC)
CELL AND CELLULAR COMPONENT	Cell (CL)	Cellular Component (CC)	Cellular Function (CC)	
MOLECULE	Molecule (CHEBI, SO, INAC, FBC)	Molecular Function (CC)	Molecular Process (CC)	

Fig. 2. The structure of the OBO Foundry (shaded regions correspond to the three original GO ontologies).



Fig. 3. Sample pages from Woodger's *Axiomatic Method in Biology* (1937).

Table 1

List of primitive classes and relations in Woodger's axiomatization of biology

Primitive classes	Primitive relations
Cell	Part of
Male gamete	Earlier than
Female gamete	Derives by division or fusion from
Whole organism	Environment of
Organized unity	
Genetic property	