

Beyond Concepts: Ontology as Reality Representation

Barry Smith

Department of Philosophy, University at Buffalo, NY 14260, USA
Institute for Formal Ontology and Medical Information Science,
Saarland University, 66041 Saarbrücken, Germany

Abstract. There is an assumption commonly embraced by ontological engineers, an assumption which has its roots in the discipline of knowledge representation, to the effect that it is *concepts* which form the subject-matter of ontology. The term ‘concept’ is hereby rarely precisely defined, and the intended role of concepts within ontology is itself subject to a variety of conflicting (and sometimes intrinsically incoherent) interpretations. It seems, however, to be widely accepted that concepts are in some sense the products of human cognition.

The present essay is devoted to the application of ontology in support of research in the natural sciences. It defends the thesis that ontologies developed for such purposes should be understood as having as their subject matter, not concepts, but rather the universals and particulars which exist in reality and are captured in scientific laws. We outline the benefits of a view along these lines by showing how it yields rigorous formal definitions of the foundational relations used in many influential ontologies, illustrating our results by reference to examples drawn from the domain of the life sciences.

1 Idealism

It is a matter of considerable astonishment to ontology-minded philosophers that many thoughtful members of the knowledge representation and related communities, including many of those involved in the development of ontologies, have embraced one or other form of idealist, skeptical, or constructionist philosophy. This means for example:

- a) a view according to which there is no such thing as objective reality to which the concepts or general terms in our knowledge representation systems would correspond;
- b) a view according to which we cannot know what objective reality is like, so that there is no practical benefit to be gained from the attempt to establish such a correspondence;
- c) a view according to which the term ‘reality’ in any case signifies nothing more than a construction built out of concepts, so that every concept-system would in principle have an equal claim to constituting its own ‘reality’ or ‘possible world’.

Doctrines under all three headings nowadays appear commonly in the wider world under the guise of postmodernism or cultural relativism, where they amount to a thesis according to which the theories of objective reality developed by the natural sciences are nothing more than cultural constructs, comparable to astrology or witchcraft. In the AI world they are often associated with constructivist ideas, for example as propounded by Maturana [1],

who holds that even biology and physics do not reflect any objective reality but are designed, rather, to help us adapt to a world which we ourselves create through our subjective experiences.

Gruber notoriously defines 'ontology' as 'a specification of a conceptualization' [2], and definitions in Gruberian spirit have been and still are accepted by most ontological engineers. A recent example is provided by the website owlseek.com, which provides the following definition of ontology:

We can never know reality in its purest form; we can only interpret it through our senses and experiences. Therefore, everyone has their own perspective of reality. An ontology is a formal specification of a perspective. If two people agree to use the same ontology when communicating, then there should be no ambiguity in the communication. To enable this, an ontology codifies the semantics used to represent and reason with a body of knowledge.

Such views have become entrenched not least because much work in ontology rests on practices predominant in the field of knowledge representation, where it is assumed as a matter of course that knowledge representation has to do not with reality but rather with concepts conceived as human creations.

A first argument for this assumption might be formulated as follows. Knowledge exists in the minds of human subjects. Hence we can have knowledge of things in reality only insofar as they are brought under the conditions which are the presuppositions of their being taken up into our minds. Hence we can have knowledge not of entities as they are in themselves but only of our own concepts.

As David Stove points out, this argument has the same form as:

We can eat oysters only insofar as they are brought under the physiological and chemical conditions which are the presuppositions of the possibility of being eaten. Hence we cannot eat oysters as they are in themselves. [3]

A second argument starts out from the premise that what we now know to be errors were in the past counted as belonging to knowledge. There are certainly among our current beliefs some that are misclassified in just this way. Hence knowledge must be allowed to comprehend also false beliefs, including false beliefs expressed by means of sentences involving general terms (such as 'phlogiston' or 'ether') which refer to nothing in reality but rather only to our own concepts.

The fallacious character of this argument rests on its failure to take careful account of time. Certainly we know that false beliefs were once erroneously counted as belonging to knowledge. But this does not prove that knowledge once comprehended false beliefs. Rather it shows only that false beliefs were sometimes erroneously misclassified as knowledge. Certainly part of what we currently count as knowledge is also mistakenly so classified. Yet the striking progress of science and technology (and its prodigiously cumulative character) gives us every reason to believe that the broad mass of what we count as knowledge today is classified correctly and that 'ether' and 'phlogiston' represent exceptions rather than the general rule. The appropriate response to the problem of error is thus to correct our errors as we find them, and this is so whether we are building ontologies or are engaged in any other type of scientific endeavor.

The response to the problem of empty general terms from the side of the concept-centered view has been, in contrast, to guarantee that every term has a referent effectively by insisting that all general terms refer in any case only to our concepts. Thus they abandon the goal of coming to grips with reality and substitute instead the much more easily attained goal of grasping conceptual entities that we ourselves have created.

In many contexts, of course, ontologists still deal with concepts, correctly, as analogous to, though as more abstract than, the linguistic expressions with which they are associated. Thus they talk of 'defining' concepts and of 'mapping' the concepts of different on-

tologies – understanding concepts effectively as tools (analogous to telescopes or microscopes) which we can use in order to gain cognitive access to corresponding entities in reality. Too often, however, there occurs an insidious shift in focus: concepts themselves become the very subject-matter of ontology. The ontologist’s tools become transformed into the very focus of his inquiries.

The influence of the concept-centered view is a product not merely of the roots of information systems ontology in the field of knowledge representation. It has become entrenched also in virtue of the fact that much work of ontology has been concerned with representations of domains, such as commerce, law, or public administration, where we are dealing with the products of human convention and agreement – and thus with entities which are in some sense merely ‘conceptual’. [4]

Today, however, we are facing a situation where ontologies are increasingly being developed in close cooperation with those working at the interface between the informatics disciplines and the empirical sciences, and under these conditions the concept-centered view is exerting a damaging influence on the progress of ontology. In what follows I present an analysis of the view and a sampling of some the problems which it brings in its wake. I then sketch an alternative view of ontology as a discipline rooted in the representation of universals and particulars in reality.

2 On Defining ‘Concept’

That there are few convincing attempts to define the term ‘concept’ (and related terms such as ‘conceptualization’ or ‘conceptual entity’) in the current literature of ontology follows in part from the fact that these terms deal with matters so fundamental to our cognitive architecture (comparable in this respect to terms like ‘identity’ or ‘object’) that attempts to define them are characteristically marked by the feature of circularity. Such circularity is illustrated for example by the Semantic Network of the Unified Medical Language System (UMLS) [5], which defines *idea or concept* as: ‘An abstract concept, such as a social, religious or philosophical concept.’

Occasionally more elaborate definitions are offered:

Concepts, also known as classes, are used in a broad sense. They can be abstract or concrete, elementary or composite, real or fictitious. In short, a concept can be anything about which something is said, and, therefore, could also be the description of a task, function, action, strategy, reasoning process, etc. [6]

This passage illustrates the way in which, in much of the relevant literature, concepts are not clearly distinguished from either entities in reality or names or descriptions on the side of language.

Another (randomly selected) example of this same confusion is provided by the Biological Pathways Exchange Ontology (BIOPAX), which defines ‘four basic concepts’: the ‘top-level *entity* class’ and ‘three subclasses: *pathway*, *interaction* and *physicalEntity*’. It then provides for the top-level class *entity* the following definition: ‘Any concept that we will refer to as a discrete unit when describing biological pathways, e.g. a pathway, interaction or *physicalEntity*.’ [7]

The tendency to run together concepts and entities is found also in linguistics; for example in passages like the following:

we are capable of constructing conceptual worlds of arbitrary complexity involving entities and phenomena that have no direct counterpart in peripherally connected experience. Such are the worlds of dreams, stories, mythology, mathematics, predictions about the future, flights of the imagination, and linguistic theories. All of us have constructed many conceptual worlds that differ in genre, complexity, conventionality, abstractness, degree of entrenchment, and so on. *For many linguistic purposes all of these worlds are on a par with the one we distinguish as “reality”.* [8] (Emphasis added)

But where it may be acceptable for linguists' purposes to run together cells and molecules with whatever might be the ontological correlates of myths and fairy tales, such a distinction is indispensable when we embark on the development of ontologies in support of natural science.

3 The Linguistic Reading of 'Concept'

The core reading of the term 'concept' in the knowledge representation and related literatures starts out from the recognition that different terms – for example terms in different languages such as 'dog', 'chien', and 'Hund' – may have the same meaning. 'Concept' is then used in place of 'name' or 'word' as a device which allows us to abstract away from incidental syntactic differences and focus instead on those sorts of relations between terms which are important for reasoning. Sometimes the concept is explicitly identified with the meaning that is shared in common by the relevant terms. Sometimes it is seen rather as something psychological – something like an idea shared in common in the minds of those who use these terms. Sometimes the concept is seen as a logical construct, for example as a 'synset' in WordNet terminology, i.e. as a set of words which can be exchanged for each other *salva veritate* in given sentential contexts. [9]

One obvious problem with the concept-centered view of ontology is that it is difficult to understand how ontologies could be evaluated on its basis. Intuitively, a good ontology is one which corresponds to reality as it exists beyond our concepts. If, however, knowledge itself is identified with knowledge of our concepts, and if an ontology is a mere specification of a conceptualization, then the distinction between good and bad ontologies seems to lose its foothold.

This problem would in other disciplines rightly be regarded as grievous. How, then, is the linguistic reading able to retain its grip on its adherents? Only, I believe, because this reading is rarely wielded without alien admixtures. 'Concept' is for example used in such a way that it is assumed to carry also connotations normally associated with terms such as 'property', 'kind' or 'universal' – terms which in normal usage do not denote entities which are the products of human cognition. It is this additional baggage which is responsible for the preponderance of confused interpretations of 'concept' adverted to already above.

Could we, then, reform the literature of knowledge representation by enforcing the linguistic reading in consistent fashion? Unfortunately not. For this would have the effect of collapsing KR into a branch of psychology or anthropology, a discipline whose claim to the effect that it is modeling the *knowledge* possessed by human subjects – rather than their mere *beliefs* – would lose its force. For uses of language to express both true and false beliefs are, after all, from the linguistic perspective cut of one cloth. One is not capturing knowledge when one describes the beliefs widely distributed in certain cultures pertaining to concepts such as *alien implant removal* or *Chios energy healing*.

4 *Is_a* and the Linguistic Reading

On the linguistic reading the assertions of *is_a* and other relations between concepts are assertions about meanings or ideas. A sentence like

lytic vacuole is_a vacuole

is, appearances notwithstanding, not an assertion about lytic vacuoles; rather it is an assertion about language use. It tells us that the meaning associated with the name 'lytic vacuole' is *narrower* or *more specific* than the meaning associated with the name 'vacuole' by this or that group of subjects.

This interpretation is, as we would anticipate, especially common in work on terminologies and thesauri. There we are interested not in *is_a* relations in the strict sense (and not in scientific laws), but rather only in various kinds of relations of ‘association’ between concepts and in the networks which these form. Statistics-based pattern recognition techniques can be applied to such networks in support of a range of information retrieval and extraction tasks – and for these purposes it may be of no account that the concepts of which such networks are formed fail to correspond to any external reality.

Note that matters are not essentially changed when the linguistic reading is given a precise technical specification, as for example in the Standard Upper Ontology, which defines a *SUO_concept* as a tuple $(p, t, d, [s])$, in which

p is a predicate defined by a definition or axioms in KIF;

t is an English term (word or multiword phrase);

d is an English documentation which attempts to precisely define the term;

s is an optional English syntactic category represented by one of the following character strings:

“noun”, “intransitive verb”, “transitive verb”, [etc.] [10]

Here, too, sentences like ‘*lytic vacuole is_a vacuole*’ turn out to be transformed into sentences which are not about vacuoles at all, but rather (somewhat implausibly) about set-theoretic objects built out of syntactic strings as urelements.

Something similar applies also in the context of Description Logic (DL), where ‘concept’ is standardly used as an abbreviation for ‘concept description’ (which – again somewhat confusingly – is not in its turn to be interpreted as meaning ‘description of a concept’). [11] This means that in DL circles talk of concepts is talk of certain *syntactic* entities. Such talk is, to be sure, semantically motivated (in the sense of ‘semantics’ that we know from Tarski and from set-theoretic model theory). Each DL concept description represents, with respect to any given interpretation, a collection of objects that are postulated as sharing the property that is specified by the description. Even this does not provide an anchor for concepts in external reality, however, for the objects in question may be (and standardly are) merely abstract mathematical postulates. Thus when it is said that DL provides the terms we use in ontologies with a ‘precise semantics’, then we should bear in mind that the sense of ‘semantics’ at issue here involves recourse to a mathematical abstraction that is far removed from our normal understanding of semantics as relating to the interplay between terms, meanings and corresponding entities in reality.

5 The Engineering Reading of ‘Concept’

The difficulties with the linguistic reading have led to the crystallization of a second, engineering reading of ‘concept’, a reading best exemplified in the use of the term ‘conceptual model’.

Sometimes conceptual modelers, even conceptual modelers developing applications in support of research in the natural sciences, talk as though their business were modeling *data* or *information*. [12] Were this to be the case, then the models themselves would be at one remove from the underlying reality with which the scientists themselves have to deal. Against this, however, we can note that the term ‘information’ – like the term ‘model’ (and like the term ‘semantics’) – is itself subject to many of the same confusions that have become associated with the term ‘concept’ in the KR and related literatures. Closer inspection of actual modeling practice then reveals that the modelers in question are in fact concerned with building models of entities in reality, thus for example with building models of the organization of the genome and not just of *information* about this organization contained in this or that database.

The term ‘concept’ itself, on the engineering reading, refers to entities that are created by modelers. Concepts are creatures of the computational realm which exist (in some sense hard to explain) through their representations in software, in UML diagrams, XML representations, in systems of axioms, or what one will. Such creation of concepts need not be a trivial matter. Not every collection of lines of code is interpretable as being associated with a conceptual model. To count as thus interpretable the code must pass what we might call a simulation test. This means that the code on execution must be such that there are relations between inputs and outputs which match relations between corresponding entities in reality. The relevant input- and output-concepts as joined together functionally by the software must in this sense *model*, which means (we presume): stand in some sort of isomorphism to, corresponding entities in reality. To the extent that the engineering reading of concepts makes sense at all, therefore, it, too, must make appeal to reality as its exists outside our minds.

6 An Ontological Turn

Good ontology and good modeling in support of the natural sciences can, we conclude, be advanced by the cultivation of a discipline that is devoted precisely to the representation of entities as they exist in reality.

In the framework of such a discipline – which would look very much like the discipline of ontology as practiced by philosophers such as Aristotle [13], Ingarden [14], Chisholm [15], Johansson [16] or Lowe [17] – we would talk not of *concepts* as linguistic or computer artefacts but rather of *universals*, conceived as that in reality to which the general terms used in making scientific assertions correspond. The particulars or tokens with which we have to deal, for example when carrying out experiments in natural science, are the instances of such universals which exist in the real world of space and time. The term ‘universal’ then signifies what the corresponding instances – for instance all *whales*, all *enzymes* – have in common. Universals are invariants in reality.

Universals and their instances enjoy hereby a symbiotic relationship: the one cannot exist without the other. Statements like:

human is_a mammal
metabolism is_a physiological process
nucleus part_of cell
cell part_of eukaryotic organism

can be interpreted both as statements about universals and as abbreviated versions of statements about the corresponding instances.

Statements like ‘*whale is_a mammal*’ or ‘*regulation of protein kinase activity part_of protein amino acid phosphorylation*’ convey knowledge precisely because they represent relations between entities in reality, relations to which the advance of science has given us cognitive access but which themselves obtain independently of our cognitive activities. They convey not extensional relations analogous to that of set-theoretic inclusion, but rather law-like relations between universals of the sort that are discovered through scientific research. [18]

Taking the reality of universals into account also gives us a means of coming to grips with what constitutes the difference between good and bad ontologies:

Bad ontologies are (*inter alia*) those whose general terms lack the relation to corresponding universals in reality, and thereby also to corresponding instances.

Good ontologies are *reality representations*, and the fact that such representations are possible is shown by the fact that, as is documented in our scientific textbooks, very many of them have already been achieved, though of course always only at some specific level of granularity and to some specific degree of precision, detail and completeness. [19]

It seems, now, to be a presupposition of much work in the field of knowledge representation that something changes in this respect when we bring computers into play (when scientific texts are supplemented by data stored electronically) – as if terms stored in computers were for some reason incapable of relating to entities in reality in the same way as do terms printed in scientific texts. When challenged to defend their assumption that computer representations must be representations of special artefacts (‘concepts’, ‘models’, ‘strings’), the adherents of such views – if I am understanding them correctly – defend themselves by pointing out that real physical entities (cells, organisms, diseases) *cannot be stored inside the computer*. This, however, is to reveal a simple misunderstanding of the nature of representation. It is to make a mistake that is equal and opposite to that made by the academicians of Lagago, who held that

since Words are only Names for Things, it would be more convenient for all Men to carry about them, such Things as were necessary to express the particular Business they are to discourse on ... which hath only this Inconvenience attending it, that if a Man’s Business be very great, and of various kinds, he must be obliged in Proportion to carry a greater bundle of Things upon his Back. [20]

7 *Is_a* as a Relation between Universals

A view of ontology as reality representation does not merely resolve certain broadly philosophical confusions common on the part of information systems ontologists. It can also provide a quite specific kind of practical help to those involved in building ontologies in support of empirical science by offering the resources to provide formally rigorous definitions of basic ontological relations.

Universals, we said, are instantiated by instances in space and time. A formal theory of this instantiation relation is advanced in [21], [22] and [23]. It begins by drawing a primitive distinction, within the realm of entities in general, between particulars and non-particulars. Examples of particulars (individuals, tokens) are: you and me, the planet Earth, this piece of cheese, your eating of this piece of cheese. Examples of non-particulars are universals such as: *human being*, *enzyme*, *butterfly*, *heptolysis*, *death*. Universals (kinds, types) thus come in different categories. Above all we can distinguish universals instantiated by continuant objects (things, substances and their parts), and universals instantiated by occurrent processes (activities, events). We shall confine ourselves in what follows to the former.

A universal is defined as anything that is instantiated, and an instance as anything that instantiates some universal. The relation of instantiation is hereby taken as primitive, and it is specified axiomatically that it holds exclusively between instances and universals (in that order). Note that it is not the case that all particulars are instances. This is because there are what we might call junk particulars (for example the mereological sum of Bush’s right knee and the pain in Clinton’s left leg) which instantiate no universal (or at least no universal standing to other universals in relations captured by scientific laws).

We can now apply these ideas to the formal treatment of the *is_a* relation as this is used for example in biomedical ontologies such as the Foundational Model of Anatomy [24], the Gene Ontology [25], and other ontologies curated under the auspices of the Open Biological Ontologies consortium [26].

Standard definitions of ‘*is_a*’ conceive it in broadly set-theoretic terms: ‘*A is_a B*’ is held to mean that the set of instances of *A* is a subset of the set of instances of *B*. The problem with this definition is, first of all, that it makes it difficult for us to do justice to the temporal complexities of the relation between instances and universals and thus to take care of false positives such as ‘*adult is_a child*’. It makes it difficult also to take care of situations (for example in the area of embryology) where one and the same individual – for example as larva and as butterfly – may be held to instantiate different universals at different times. A second problem with the set-theoretic reading is that it admits of purely contingent cases of class subsumption, as illustrated in:

dog owned by me is_a mammal weighing less than 200 Kg
dog in Leipzig is_a dog

and also of logically constructed cases of class subsumption, such as:

dog is_a dog or apple
dog and mammal is_a dog or apple
dog is_a non-cat

Cases of these sorts are surely not admissible as genuine *is_a* relations where we are attempting to develop ontologies for purposes of supporting inquiry in disciplines like the life sciences. For the task of such ontologies is to capture not contingent classifications (effected, for example, for administrative purposes) but rather scientific laws.

We can, however, by calling in aid the theory of universals and instantiation, formulate a better definition of the ‘*is_a*’ relation which will exclude such spurious cases:

A is_a B if and only if: (1) *A* and *B* are universals, and (2) for all times *t*, if anything instantiates universal *A* at *t* then that same thing must instantiate also the universal *B* at *t*.

The phrase ‘must instantiate’ here connotes: *in virtue of a scientific law*. Our proposed definition thus involves an essential departure not only from the concept-centered view but also from the extensionalism of set-theory-based approaches, and thus also from the familiar DL-based formalisms founded thereon.

8 *Part_of* as a Relation between Universals

The power of the universals-based methodology becomes still clearer when we turn to the relation *part_of*, another relation widely used in current biomedical ontologies. For this relation can be given no coherent interpretation at all when considered as a relation between concepts on either the linguistic or the engineering reading. What would it mean, after all, to say that ‘*coccyx part_of vertebral column*’ expresses a relation of the same type as the *is_narrower_than* relation conceived as a relation between meanings? And what would it mean to say that this statement expresses a relation between artefacts of a computer model?

To understand ‘*A part_of B*’ we need first of all to call in aid the familiar mereological parthood relation understood as a relation among particulars and illustrated by: ‘Jane’s heart is part of Jane’s body’. We can then establish that assertions of the form ‘*A part_of B*’ in fact signify a variety of different sorts of relations among the instances of the corresponding universals – relations which are often confusedly run together when understood as relations between concepts. [27] Thus (and least interestingly for scientific purposes) it can mean:

some instances of *A* are parts of some instances of *B*,

as in *roof part_of car* (WordNet), or *acquired abnormality part_of bird* (UMLS). More sig-

nificantly, for the purposes of ontologies designed to support the needs of scientific research, *A part_of B* can express relations between universals such as:

A part_for B, which asserts (i) that if an instance of *A* exists at a given time, then an instance of *B* exists at this same time and (ii) that the former is an instance-level part of the latter. *A part_for B* thus provides information primarily about the *As*; it tells us that *As* do not exist except as instance-level parts of *Bs*. Examples: *nucleus part_for cell*; *human testis part_for human being*.

B has_part A, which asserts (i) that if an instance of *B* exists at a given time, then an instance of *A* exists at this same time and (ii) the latter is an instance-level part of the former. *B has_part A* thus provides information primarily about the *Bs*; it tells us that *Bs* do not exist except with *As* as instance-level parts. Examples: *woman has_part heart*; *cell has_part membrane*.

A reading of *part_of* as a relation between universals which suggests itself against this background would see this relation as a combination of *part_for* with *has_part*. Thus:

A part_of B obtains if and only if, (i) for any instance *x* of *A* existing at some time *t*, there is some simultaneously existing instance *y* of *B* which is such that *x* stands to *y* in the instance-level part relation, and (ii) *vice versa* for any instance *y* of *B*.

This implies a strong structural tie between the universals *A* and *B*: the one can be instantiated only hand in hand with the other. Examples: *cell membrane part_of cell*; *human brain part_of human being*.

The resultant framework gives us the resources to formulate definitions also of time-dependent part-relations between universals for example of the type expressed in a statement to the effect that a notochord is part of a mammal only in the embryonic state. Thus it allows us also to add improvements to existing lexical resources in other fields, for example by adding to WordNet the facility to deal with what one might call ‘optional’ body parts such as warts and freckles, and with ‘temporary optional’ body parts such as fetuses or five-o’clock shadow.

9 Relations in the UMLS Semantic Network

The framework outlined above illustrates the degree to which, when we move beyond the domain of concepts to the realm of universals and particulars in a changing reality, formal distinctions become apparent which have been too often glossed over. It also provides us with a guiding thread as to the approach that needs to be adopted to provide formally rigorous definitions also of those other relations – such as ‘*causes*’, ‘*is_a_realization_of*’, ‘*participates_in*’, ‘*develops_from*’, ‘*derives_from*’, and so on – which are standardly employed in the construction of ontologies in the life sciences. To see why a coherent formal treatment of such foundational relations is needed if ontologies are to be developed which are capable of supporting scientific inquiry in an area such as biomedicine, consider the treatment which they currently receive in the UMLS Semantic Network (SN), described as an ‘upper level ontology for the biomedical domain’ [5].

SN is a graph-theoretic structure that is designed to unify the 975,354 ‘concepts’ and 2.4 million ‘concept names’ of the UMLS Metathesaurus, itself one of the most important tools of inquiry in the domain of biomedical informatics. Inspection reveals that the majority of the 54 relations contained in SN are subject to problems even more serious than those affecting traditional treatments of ‘*part_of*’.

Table 1 consists of randomly selected examples of the 6000 or so edges which form the SN conceived as graph-theoretical structure. These examples nicely illustrate the dubious assertions which arise when one adopts a purely linguistic reading of ‘concept’, and at the same time they reveal that great difficulties will be set in the way of anyone who attempts to establish what it is, in reality, to which SN’s relation-types correspond. In what sense does a mental process *precede* a genetic function? In what sense does an antibiotic *cause* an experimental model of a disease?

Table 1. *Semantic Relations from the UMLS Semantic Network, with definitions and selected examples*

<i>precedes: occurs earlier in time</i>	
<i>Cell Function precedes Cell Function</i> <i>Mental Process precedes Molecular Function</i> <i>Molecular Function precedes Mental Process</i>	<i>Mental Process precedes Genetic Function</i> <i>Experimental Model of Disease precedes Cell or Molecular Dysfunction</i>
<i>affects: produces a direct effect on. Implied here is the altering or influencing of an existing condition, state, situation, or entity</i>	
<i>Acquired Abnormality affects Fish</i> <i>Physiologic Function affects Reptile</i>	<i>Experimental Model of Disease affects Fungus</i> <i>Experimental Model of Disease affects Genetic Function</i>
<i>causes: brings about a condition or an effect. Implied here is that an agent, such as for example a pharmacologic substance or an organism, has brought about the effect</i>	
<i>Food causes Experimental Model of Disease</i> <i>Biomedical or Dental Material causes Mental or Behavioral Dysfunction</i>	<i>Manufactured Object causes Disease or Syndrome</i> <i>Vitamin causes Injury or Poisoning</i> <i>Bacterium causes Pathologic Function</i>
<i>location_of: the position, site, or region of an entity or the site of a process</i>	
<i>Gene or Genome location_of Fungus</i> <i>Fungus location_of Vitamin</i>	<i>Organization location_of Diagnostic Procedure</i> <i>Tissue location_of Mental or Behavioral Dysfunction</i>

10 From Names to Objects

Such questions are not otiose. They cut to the very heart of ontology as a tool of contemporary biomedical informatics. Consider, for example, what the molecular biologist Sydney Brenner has to say in his critical comments on the Gene Ontology, published under the title “Ontology Recapitulates Philology”. [28] Brenner first of all accepts that we ‘need a theoretical framework in which to embed biological data’ so that the endless streams of data currently being assembled through biomedical research ‘can be sifted and abstracted.’ He then goes on to insist, rightly, that ‘the network we should be interested in is not the network of *names* but the network of *the objects themselves*.’ (Emphasis added.)

At the same time, however, he asserts that it is only when exploring reality at the granularity of molecules that one has to do with objects. Only there, he tells us, do objects ‘have their own names: they are chemical names written in the language of DNA sequences’. Brenner unfortunately reveals hereby that (like the academicians of Lagago) he, too, has a very weak understanding of the relation between a name and its denotatum. At the same time he denies what the Gene Ontology correctly recognizes, namely that biologi-

cal reality can be grasped in adequate fashion only by taking entities at a multiplicity of different granularities into account.

Where the Gene Ontology goes wrong is in its failure to provide even the beginnings of a formal architecture for its *is_a* and *part_of* relations and thus also – since granular hierarchies are structured via *part_of* relations – for its treatment of the different granular levels which fall within its scope. [29] And if the argument presented above is correct, then this lack of an adequate formal treatment of relations, which is manifested also by the UMLS SN, flows precisely from a failure to penetrate beyond the realm of concepts linguistically conceived to the world of universals and their instances in spatiotemporal reality. It is thus especially poignant that Lawrence Hunter, in his response to Brenner on behalf of GO [30], asserts that ‘the essence of the Gene Ontology project ... and of other knowledge-bases of molecular biology [such as] the Unified Medical Language System, is not in the list of names they embody, but in the relationships they represent’. For as has been repeatedly pointed out, it is these very relationships which are currently dealt with so incoherently in the knowledge-bases mentioned.

As Hunter points out, ontologies are designed primarily for use not by human beings but rather by computer programs in order to accomplish complex inference tasks. This however, as he also points out, ‘requires the presence of a well represented knowledge-base of molecular biological entities.’ If neither the Gene Ontology nor the UMLS as a whole yet comprehends a well-represented knowledge-base of the needed sort, then this is at least in part because they have been constructed within an ontological framework in which the necessary coming to grips with entities in reality has been blocked by the focus on their linguistic surrogates in the realm of concepts.

Acknowledgements

Work on this paper was carried out under the auspices of the Wolfgang Paul Program of the Alexander von Humboldt Foundation, the EU Network of Excellence in Medical Informatics and Semantic Data Mining, and the Project “Forms of Life” sponsored by the Volkswagen Foundation. Thanks go also to Bill Andersen, Sebastian Brandt, James Cimino, Dirk Siebart and Leo Zaibert for helpful comments.

References

- [1] Humberto R. Maturana and Francisco J. Varela, *The Tree of Knowledge: The Biological Roots of Human Understanding*, 1998, Shambhala Publications, Boston MA
- [2] T. R. Gruber, A Translation Approach to Portable Ontologies, *Knowledge Acquisition*, 5: 199-220, 1993.
- [3] James Franklin, Stove’s Discovery of the Worst Argument in the World, *Philosophy*, 77: 615-24, 2002.
- [4] Andrew Frank, Ontology: A Consumer’s Point of View, *Spatial and Temporal Reasoning*, ed. Oliviero Stock, Kluwer, Dordrecht, 1996.
- [5] Alexa T. McCray, Representing biomedical knowledge in the UMLS semantic network, *High performance medical libraries: Advances in information management for the virtual era*, 45-55, 1993.
- [6] Oscar Corcho and Asuncion Gomez-Perez, A Roadmap to Ontology Specification Languages, in Rose Dieng and Olivier Corby (eds.), *Knowledge Engineering and Knowledge Management. Methods, Models and Tools*, Springer, Berlin, 80-96, 2000.
- [7] http://www.biopax.org/Downloads/Level1v0.5.2/biopax-level1-v0.5.2_Ontology_Documentation.pdf.
- [8] Ronald Langacker, *Foundations of Cognitive Grammar*, Stanford University Press, 1987/1991.
- [9] Christiane Fellbaum, ed., *WordNet: An Electronic Lexical Database*, Cambridge, MA: MIT Press, 1998.
- [10] <http://suo.ieee.org/email/msg01175.html>
- [11] Franz Baader, Diego Calvanese, Deborah McGuinness, Daniele Nardi, and Peter Patel-Schneider, Editors, *The Description Logic Handbook: Theory, Implementation and Applications*, Cambridge Univer-

- sity Press, Cambridge, 2003.
- [12] Norman S. Paton, *et al.*, Conceptual Modelling of Genomic Information, *Bioinformatics*, 16: 6, 548–557.
 - [13] S. Marc Cohen, Aristotle’s Metaphysics, *Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/entries/aristotle-metaphysics>.
 - [14] Roman Ingarden, *Der Streit um die Existenz der Welt*, Tübingen: Niemeyer, 3 volumes, 1964/1965/1974.
 - [15] Roderick M. Chisholm, *A Realistic Theory of Categories: An Essay on Ontology*, Cambridge: Cambridge University Press, 1996.
 - [16] Ingvar Johansson, *Ontological Investigations. An Inquiry into the Categories of Nature, Man and Society*, Frankfurt: Ontos, 2004.
 - [17] E. J. Lowe, *A Survey of Metaphysics*, Oxford: Oxford University Press, 2002.
 - [18] David M. Armstrong, *Universals and Scientific Realism*, Volume 1: *Nominalism and Realism*, Volume 2: *A Theory of Universals*, Cambridge: Cambridge University Press, 1978.
 - [19] Thomas Bittner and Barry Smith, A Theory of Granular Partitions, in: M. Duckham, *et al.* (eds.), *Foundations of Geographic Information Science*, Taylor & Francis, London, 117-151, 2003.
 - [20] Jonathan Swift, *Gulliver’s Travels*, Part III, A Voyage to Laputa, Balnibarbi, Luggnagg, Glubbdubdrin and Japan, Chapter V, 1726/1735.
 - [21] Jan Berg, Aristotle’s Theory of Definition, *ATTI del Convegno Internazionale di Storia della Logica*, San Gimignano, 4–8 December 1982, Bologna: CLUEB, 1983, 19–30.
 - [22] Barry Smith, The Logic of Biological Classification and the Foundations of Biomedical Ontology, in: Dag Westerstahl (ed.): *Invited Papers from the 10th International Conference in Logic, Methodology and Philosophy of Science*, Oviedo, Spain, 2003.
 - [23] Fabian Neuhaus, Pierre Grenon and Barry Smith, A Formal Theory of Substances, Qualities, and Universals, *in this volume*.
 - [24] C. Rosse and J. L. V. Mejino, A Reference Ontology for Bioinformatics: The Foundational Model of Anatomy. *Journal of Biomedical Informatics*, 36: 478-500, 2003.
 - [25] <http://www.geneontology.org/>
 - [26] <http://obo.sourceforge.net/>
 - [27] Barry Smith and Cornelius Rosse, The Role of Foundational Relations in the Alignment of Biomedical Ontologies, *Proceedings of Medinfo*, 2004.
 - [28] Sydney Brenner Life Sentences: Ontology Recapitulates Philology, *Genome Biology* 3 (4):1006.1-1006.2, 2002.
 - [29] Barry Smith, Jacob Köhler and Anand Kumar, On the Application of Formal Principles to Life Science Data: A Case Study in the Gene Ontology, *Proceedings of DILS 2004. Data Integration in the Life Sciences* (Lecture Notes in Computer Science 2994). Springer, Berlin, 124-139, 2004.
 - [30] Lawrence Hunter, Ontologies for Programs, Not People, *Genome Biology* 3 (6): 1002.1–1002.2, 2002.