

The Logic of Biological Classification and the Foundations of Biomedical Ontology

Barry Smith

Department of Philosophy, University at Buffalo and Institute for Formal Ontology and Medical Information Science, University at Leipzig

To appear in Dag Westerståhl (ed.), *Invited Papers from the 10th International Conference in Logic Methodology and Philosophy of Science*, Oviedo, Spain, 2003 (Elsevier-North-Holland, 2004)

Abstract

Biomedical research is increasingly a matter of the navigation through large computerized information resources deriving from functional genomics or from the biochemistry of disease pathways. To make such navigation possible, controlled vocabularies are needed in terms of which data from different sources can be unified. One of the most influential developments in this regard is the so-called Gene Ontology, which consists of controlled vocabularies of terms used by biologists to describe cellular constituents, biological processes and molecular functions, organized into hierarchies via the relation of class subsumption. Here we seek to provide a rigorous account of the logic of classification that underlies GO and similar biomedical ontologies. Drawing on Aristotle, we develop a system of axioms and definitions for the treatment of biological classes and instances.

Introduction

In reflection of the huge amounts of data accumulating in areas such as genomics and proteomics, biology and biomedicine have come to rely increasingly on the use of computational methods in their research. One of the most impressive and influential developments in this regard is the so-called Gene Ontology (GO),¹ which is being developed as part of the effort to produce controlled vocabularies for shared use across different biological domains within the framework of the Open Biological Ontologies project.² We take GO as our test case in what follows, not only because it has proved so successful in serving as a common reference system for a variety of groups working at the forefront of biomedical research, but also because, as we shall see, it suffers from a series of problems which are characteristic of almost all current ontologies used in bioinformatics.

¹ The Gene Ontology Consortium, "Gene Ontology: Tool for the Unification of Biology. *Nature Genetics*, 25 (2000), 25-29. See also: <http://www.geneontology.org>.

² <http://obo.sourceforge.net>.

GO provides some 20,000 terms for describing gene product attributes. It is divided into three hierarchically structured networks, whose topmost nodes are, respectively: *cellular component*, *molecular function* and *biological process*.³ While GO is not strictly speaking an ontology in the sense in which this term is understood by philosophers, it does go some way in this direction, in that its three constituent vocabularies are organized as hierarchies via the ontological relations of subsumption (*human being* is subsumed by *mammal*) and partonomic inclusion (*human heart* is included as part of *human being*). Following standard usage in GO and other similar endeavors, these relations are called ‘*is a*’ and ‘*part of*’ in what follows.

Here we are concerned with GO as a *classification* of biological phenomena. The classes which stand in its *is a* and *part of* relations have some obvious relation to the species and genera of more traditional biological classifications, but there are also important differences. Thus not only are classes of *objects* recognized by GO, but so too are classes of *processes* and *functions*.⁴ Crucially, GO defines its three structured networks as separate ontologies, which means that no ontological relations are defined between them. In other respects, too, the GO literature provides few clues as to how the ontological correlates of its separate constituent terms are to be conceived. Thus in particular, it tells us little about how we are to understand the two central terms *biological process* and *molecular function*.⁵

As a step towards filling this gap, and in reflection of the fact that GO, like many other ontologies currently being developed for purposes of biomedical research, shuns logico-philosophical rigor, we provide here a formal account of biological and biomedical classification which is designed as a first step towards the rigorous treatment of the questions concerning classes and class-hierarchies which arise at the interface between biology and medicine on the one hand and current bioinformatics research on the other.

The Gene Ontology

For purposes of preliminary orientation, consider the two GO terms:

GO:0003673: cell fate commitment

GO:0045168: cell-cell signaling involved in cell fate commitment

The hierarchical relations between these two entries within GO’s biological process ontology are shown in Figure 1 below.

‘Is a’, as it is employed in this diagram, means roughly what we would expect it to mean when interpreted as a relation of subsumption between classes (natural kinds, species, genera) in biology. Note, though, that (unlike Aristotle, and unlike Linnaeus) GO

³ <http://www.geneontology.org/doc/GO.doc.html>. We refer in what follows to the version of October 2003.

⁴ On the different logical frameworks needed for the treatment of objects, functions and processes see Pierre Grenon and Barry Smith, “SNAP and SPAN: Towards Dynamic Spatial Ontology”, forthcoming in *Spatial Cognition and Computation*.

⁵ See Barry Smith, Jennifer Williams and Steffen Schulze-Kremer, 2003, “The Ontology of the Gene Ontology”, in *Biomedical and Health Informatics: From Foundations to Applications*, Proceedings of the Annual Symposium of the American Medical Informatics Association, Washington DC, November 2003, 609–613.

allows multiple inheritance; that is to say, it allows one and the same biological class to have two or more parent-classes (as, in the figure, *cell differentiation* has the two parents *development* and *cellular process*). In addition GO does not strive to ensure that the terms in its three hierarchies are divided into predetermined levels (analogous to the levels of kingdom, phylum, class, order, etc., in traditional biology); indeed the acceptance of multiple inheritance means that such levels cannot in any case be defined, since the notion of ‘sibling’ becomes indeterminate.

Multiple inheritance allows us to deal with different aspects and contexts of classification within a single network. It is thus a useful device for producing compact networks which can facilitate computationally efficient navigation through large edifices of information.

At the same time, however, multiple inheritance causes problems. These turn *inter alia* on the fact that the alignment of distinct ontologies rests crucially on the assumption that the basic ontological relations – above all relations such as *is a* and *part of*, which provide the glue which holds ontologies together – must have the same meanings in the different ontologies to be aligned. As inspection reveals, however, multiple inheritance goes hand in hand, at least in many cases, with the assignment to the *is a* relation of a variety of meanings within a single ontology. The resultant *mélange* makes coherent integration across ontologies achievable (at best) only under the guidance of human beings with the sorts of biological knowledge which can override the mismatches which otherwise threaten to arise. This, however, is to defeat the very purpose of constructing bioinformatics ontologies like GO as the basis for a new kind of biological and biomedical research designed to exploit the power of computers.⁶

Thus for example when GO postulates

cell differentiation *is a* cellular process

cell differentiation *is a* development

then it means two different things by ‘is a’. Only in the former case do we have to deal with a true subsumption relation between biological classes. In the latter case, rather, as is seen from the definition:

GO:0007275 **Development**

Definition: Biological processes specifically aimed at the progression of an organism over time from an initial condition (e.g. a zygote, or a young adult) to a later condition (e.g. a multicellular animal or an aged adult)

the relation involved would more properly be expressed as: *contributes to the achievement of a certain end*.

⁶ See Barry Smith, Jakob Köhler and Anand Kumar, “On the Application of Formal Principles to Life Science Data: A Case Study in the Gene Ontology”, in *Proceedings of DILS 2004 (Data Integration in the Life Sciences)*, (Lecture Notes in Bioinformatics 2994), Berlin: Springer, 2004.

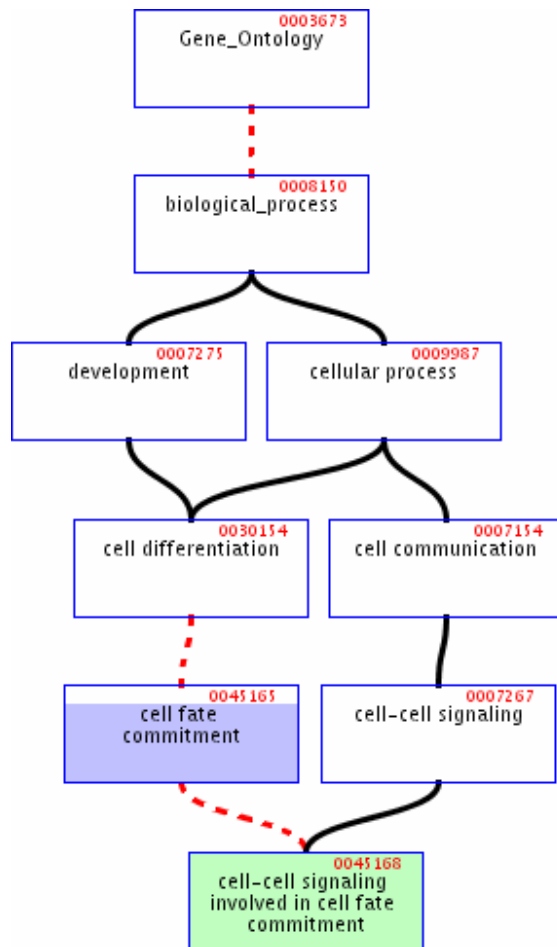


Figure 1: Example of GO Relations.⁷

When GO postulates:

hexose biosynthesis *is a* monosaccharide biosynthesis

hexose biosynthesis *is a* hexose metabolism,

on the other hand, then the second *is a* seems more properly to amount to a *part of* relation, since *hexose biosynthesis* is just that part of hexose metabolism in which hexose is synthesized.

And when GO postulates:

vacuole (sensu Fungi) *is a* storage vacuole

⁷ The diagram is taken from the QuickGO browser: <http://www.ebi.ac.uk/ego>. Solid links indicate *is a* relations; broken links indicate *part of* relations.

vacuole (sensu Fungi) *is a* lytic vacuole,

where the ‘sensu’ operator is introduced by GO to cope with those cases where a word or phrase has a specific meaning when applied to specific classes of organisms,⁸ then it seems that *is a* stands in neither case for a genuine subsumption relation between biological classes; rather, it signifies on the one hand the assignment of a *function* and on the other hand the assignment of special features to the entities in question.⁹ The case is thus analogous to:

tank (sensu Oil Industry) *is a* storage tank

tank (sensu Oil Industry) *is a* tank with an enamel coating to prevent rust.

The term ‘tank’ as used in the oil industry designates in every case a tank used for storage, and all such tanks have an enamel coating to prevent rust. But in neither case do we have what should properly be represented as an *is a* relation in a well-designed ontology.

Theorists of classification have long recognized that the division into levels and the possession by every level within a classificatory hierarchy of the so-called JEPD property (for: jointly exhaustive and pairwise disjoint) represent ideals to which classifications should aspire. The feature of exhaustivity may be difficult to achieve in the realm of biological phenomena. But shortfalls from disjointness are easy to detect. The acceptance of multiple inheritance is just the rejection of the criterion of disjointness and thus also of the JEPD ideal.

We here leave open the question whether division into levels and single inheritance involving genuine *is a* relations can be achieved throughout the realm of classifications treated of by GO and similar ontologies. However, we note that, as Guarino and Welty have shown,¹⁰ methods exist which have demonstrated considerable success in removing cases of multiple inheritance from class hierarchies by distinguishing *is a* relations from ontological relations of other sorts. Using their methods, well-structured classifications can be achieved by recognizing additional relation-types (for example: *has role*, *is dependent on*, *causes*, *is involved in*, *is realized in*) and by allowing within a single ontology categories of entities of different sorts (for instance *roles*, *functions*, *qualities*, *processes*). GO, however, has neither of these alternatives at its disposal because of its insistence that its three constituent vocabularies represent *separate ontologies* with no relations defined between them.

Core Axioms for a Theory of Biological Classification

We shall focus, in what follows, on the logical treatment of the notion of class as a step towards building a framework within which issues of biological classification can be

⁸ <http://www.geneontology.org/doc/GO.usage.html#sensu>.

⁹ A lytic vacuole is defined by GO as: a vacuole that is maintained at an acidic pH and which contains degradative enzymes, including a wide variety of acid hydrolases.

¹⁰ See e.g. their paper Nicola Guarino and Chris Welty, “Identity and subsumption”, in R. Green, C. A. Bean, and S. Hyon Myaeng (eds.), *The Semantics of Relationships: An Interdisciplinary Perspective*, Dordrecht: Kluwer (2002), 111–126.

more rigorously addressed. One might at first suppose that the logic of classes is a matter properly to be treated on a more general level – for example as part of set theory in the mathematical sense. If, however, a class is the ontological correlate of a node in a (biological) classification, if, in other words, a class is a (biological) *natural kind*, then this means that classes must stand to their instances in a relation which is quite different from the relation between a set and its members. This is because classes, but not sets, can remain identical even while undergoing a certain turnover in their instances.¹¹

Our formal theory is motivated by the theory of classes that we find in Aristotle’s writings. We turn to Aristotle not only because many of his ideas still have an astonishing pertinence when it comes to laying down standards of logical rigor in the construction of classifications and in the formulation of definitions, but also because, while many Aristotelian ideas were cast aside in the wake of the Darwinian revolution in biology, his ideas on classes and classification have in recent times come to enjoy a new relevance as a result of the role of classificatory ontologies in contemporary bioinformatics.

The theory here set forth is designed as a central module to be extended and modified to deal with specific issues relating to biological classification or with specific kinds of biological classes. As we should expect, given the Aristotelian roots of the axioms presented, the theory works well when applied to the classification of organisms and of spatially extended objects (endurants, continuants, things, substances) in general. Amended versions will be needed where we are dealing with the classification of entities, such as functions and processes, in other categories.

We begin by drawing a distinction, within the realm of entities in general, between universals and particulars. We take the opposition between universals and particulars as a primitive of our theory, and introduce variables ϵ , f , g , ... to range over entities in general. We then adopt the axiom:

$$A1. \quad \neg \exists \epsilon (u(\epsilon) \wedge p(\epsilon))$$

where u and p are primitive predicates holding of *universals* and *particulars*, respectively. Thus A1 asserts that there is nothing that is both a universal and a particular.

Examples of particulars are: you and me, the Planet Earth, this piece of cheese. Examples of universals are: *human being*, *enzyme*, *aspirin*. Particulars (individuals, tokens) are simply located entities, bound to a specific (normally topologically connected) location in space and time. Universals are multiply located entities; they exist *in* the corresponding particulars.¹²

We introduce a primitive relational predicate *inst* to stand for the relation between an instance and a class. We then define a class as anything (any universal) that is instantiated, and an instance as anything (any particular) that instantiates some class:

$$D1. \quad \text{class}(\epsilon) =_{\text{def}} \exists f \text{ inst}(f, \epsilon)$$

¹¹ See Barry Smith and Cornelius Rosse, “The Role of Foundational Relations in Biomedical Ontologies”, *Proceedings of Medinfo*, 7-11 September 2004, in press.

¹² For a formal treatment of these notions see Smith, “On Substances, Accidents and Universals: In Defence of a Constituent Ontology”, *Philosophical Papers*, 26 (1997), 105–127.

$$D2. \quad \text{instance}(e) =_{\text{def}} \exists f \text{ inst}(e, f)$$

By admitting the predicate *inst* and treating terms for classes as logically on a par with terms for instances in this way, we can develop our theory exclusively within the framework of first order logic. We might call it: first order logic with universal terms and certain designated (relational) predicates – above all identity and instantiation – which have a fixed semantic evaluation in every model.¹³

Most importantly for our purposes, the realm of universals comprehends (biological) classes, i.e. what in other contexts would be called natural kinds, species, genera, and the like. We can now postulate further:

$$A2. \quad \exists e(\mathbf{u}(e) \wedge \neg \text{class}(e))$$

There exists at least one universal which is not a class.

Examples of universals which are not classes are: *pet*, *adult*, *rational being*, *parent*, *catalyst*, *movement*, *process of development*, *storage vacuole*.¹⁴ Classes are, as it were, elite entities within the realm of universals.¹⁵ Which classes (and thus which instances) exist in a given domain is a matter for empirical research. In the macroscopic biological realm, at least, we can assume that the question as to which classes of entities exist has to do with the question as to which entities result from the coordinated expression of genes of specific sorts.

Instances, similarly, are elite entities within the realm of particulars; they are the *natural* (or standard or prototypical or canonical) exemplars of biological classes. The problems raised by non-standard instances must be dealt with in the extended version of the core module here presented, as also must the problems raised by non-standard classes, by classes in non-standard situations (for example organism species on the verge of extinction) and by the ways in which biological classes can change (evolve) over time.¹⁶

We need an axiom to the effect that:

$$A3. \quad \forall e \forall e'(\text{inst}(e, e') \rightarrow \mathbf{p}(e) \wedge \mathbf{u}(e'))$$

We can then prove the theorems:

$$T1. \quad \forall e(\text{class}(e) \rightarrow \mathbf{u}(e))$$

¹³ An alternative approach, which embraces a second-order logical framework, is explored in Nino Cocchiarella, “On the Logic of Natural Kinds”, *Philosophy of Science*, 1976; 43: 202–222.

¹⁴ In an alternative formulation of these ideas we might distinguish different *contexts of classification*. It might then be that in certain special contexts of inquiry some of these terms can indeed be held to designate classes satisfying axioms very much like the ones presented here.

¹⁵ Our theory of classes is thus an analogue of the ‘sparse theory of universals’ propounded by David Lewis in his “New Work for a Theory of Universals”, *Australasian Journal of Philosophy*, 61 (1983): 343–377 and embraced, *inter alia*, by David Armstrong.

¹⁶ Some indications are provided in Thomas Bittner and Barry Smith, “A Theory of Granular Partitions”, *Foundations of Geographic Information Science*, Matthew Duckham, Michael F. Goodchild and Michael F. Worboys, eds., London: Taylor & Francis, 2003, 117–151.

There are no classes which are not universals.

$$T2. \quad \forall \epsilon (\text{instance}(\epsilon) \rightarrow \mathbf{p}(\epsilon))$$

There are no instances which are not particulars.

$$A4. \quad \exists \epsilon (\mathbf{p}(\epsilon) \wedge \neg \text{instance}(\epsilon))$$

There are particulars which are not instances.

$$A5. \quad \exists \epsilon \mathbf{p}(\epsilon) \rightarrow \exists \epsilon' \text{instance}(\epsilon')$$

If there is a particular, then there is an instance.

As an example of a particular which is not an instance consider the mereological sum of a molecular at the end of your nose and your brother's lizard. Intuitively, every particular is such as to overlap mereologically with some instance.

We can then prove:

$$T3. \quad \neg \exists \epsilon (\text{class}(\epsilon) \wedge \text{instance}(\epsilon))$$

Nothing can be both an instance and a class.

$$T4. \quad \forall \epsilon (\text{class}(\epsilon) \rightarrow \exists \epsilon' (\mathbf{inst}(\epsilon', \epsilon)))$$

Every class has at least one instance. (This follows trivially from D1.) This is the basic principle of Aristotelian realism as far as classes are concerned. We here leave open whether an analogous axiom holds for universals in general.

$$T5. \quad \exists \epsilon \mathbf{p}(\epsilon)$$

$$T6. \quad \exists \epsilon \text{instance}(\epsilon)$$

$$T7. \quad \exists \epsilon \text{class}(\epsilon)$$

$$T8. \quad \exists \epsilon \mathbf{u}(\epsilon)$$

There exists at least one particular; there exists at least one instance; there exists at least one class; and there exists at least one universal.

We can now introduce typed variables, A, B, C, \dots to range over classes and x, y, z, \dots to range over instances, and we can postulate an axiom to the effect that at least two classes exist:

$$A6. \quad \exists A \exists B (A \neq B),$$

together with an axiom of extensionality:

$$A7. \quad \forall A \forall B \forall x ((inst(x, A) \leftrightarrow inst(x, B)) \rightarrow A = B).$$

(We note that the relation to time must be taken into account in the extended version of the core module here presented. We should then, for example, be able to formulate principles to the effect that classes are identical if and only if they share the same instances *at the same times*.)

We can now define the *is a* relation between classes in terms of *inst*:

$$D3. \quad A \text{ is a } B =_{\text{def}} \forall x (inst(x, A) \rightarrow inst(x, B)).$$

Is a is thus superficially analogous to the usual set-theoretic subset relation (\subseteq). More perspicuously:

$$D3^* \quad \epsilon \text{ is a } f =_{\text{def}} \text{class}(\epsilon) \wedge \text{class}(f) \wedge \forall x (inst(x, \epsilon) \rightarrow inst(x, f)).$$

We can also define various predicates picking out special sorts of classes, as follows:

$$D4. \quad \text{genus}(A) =_{\text{def}} \text{class}(A) \wedge \exists B (B \text{ is a } A \wedge B \neq A)$$

$$D5. \quad \text{species}(A) =_{\text{def}} \text{class}(A) \wedge \exists B (A \text{ is a } B \wedge B \neq A)$$

$$D6. \quad \text{lowestspecies}(A) =_{\text{def}} \text{species}(A) \wedge \neg \text{genus}(A)$$

$$D7. \quad \text{highestgenus}(A) =_{\text{def}} \text{genus}(A) \wedge \neg \text{species}(A)$$

Aristotle uses the term ‘category’ as a synonym of highest genus, and we can guarantee axiomatically that at least one such highest genus exists:

$$A8. \quad \exists A \text{ highestgenus}(A)$$

Adding:

$$A9. \quad \text{class}(A) \rightarrow \text{genus}(A) \vee \text{species}(A)$$

we can then prove:

$$T9. \quad \text{class}(A) \rightarrow (\text{genus}(A) \vee \text{lowestspecies}(A))$$

$$T10. \quad \text{class}(A) \rightarrow (\text{species}(A) \vee \text{highestgenus}(A))$$

and also:

- T11. $A \text{ is a } A$ (is a is reflexive)
- T12. $(A \text{ is a } B \wedge B \text{ is a } C) \rightarrow A \text{ is a } C$ (is a is transitive)
- T13. $(A \text{ is a } B \wedge B \text{ is a } A) \rightarrow A = B$ (is a is antisymmetric)

Axioms for *Nearest Species*

When one class is immediately subsumed by another (i.e. where one is child to the other as parent in a species-genus tree) then we say that they stand in the relation of nearest species, which is defined as follows:

$$\text{D8. } \text{nearestspecies}(A, B) =_{\text{def}} A \text{ is a } B \wedge A \neq B \\ \wedge \forall C ((A \text{ is a } C \wedge C \text{ is a } B) \rightarrow (C = A \vee C = B))$$

We can now formulate a series of axioms for biological classes which seem to come close to capturing what we mean when we say that classes are *natural* kinds. Here (following Aristotle¹⁷) we focus on axioms for classes of *objects* (cells, molecules, organisms, limbs, organs, and the like), noting again that the framework will in due course need to be expanded to cope with the class-instance relations governing entities in other categories:

$$\text{A10. } (\text{nearestspecies}(A, B) \wedge \text{nearestspecies}(A, C)) \rightarrow B = C$$

A species never has two *is a* parents. (This rules out cases of multiple inheritance.)

$$\text{A11. } \text{lowestspecies}(A) \wedge \text{lowestspecies}(B) \wedge A \neq B \\ \rightarrow \neg \exists x (\text{inst}(x, A) \wedge \text{inst}(x, B))$$

Distinct lowest species never share instances.

$$\text{A12. } \text{genus}(A) \wedge \text{inst}(x, A) \rightarrow \exists B \text{ nearestspecies}(B, A) \wedge \text{inst}(x, B)$$

Every instance of a genus instantiates also some nearest species of this genus.

$$\text{A13. } \text{nearestspecies}(A, B) \rightarrow \exists x (\text{inst}(x, B) \wedge \neg \text{inst}(x, A))$$

Each genus includes more instances than any of its nearest species.

$$\text{A14. } \text{nearestspecies}(B, A) \rightarrow \exists C (\text{nearestspecies}(C, A) \wedge B \neq C)$$

Every genus has at least two children.

¹⁷ More precisely: following Jan Berg's excellent treatment of these matters in: "Aristotle's Theory of Definition", *ATTI del Convegno Internazionale di Storia della Logica*, San Gimignano, 4–8 December 1982, Bologna: CLUEB, 1983, 19–30. See <http://ontology.buffalo.edu/bio/berg.pdf>.

$$\text{A15. } (\text{nearestspecies}(B, A) \wedge \text{nearestspecies}(C, A) \wedge \exists x(\text{inst}(x, B) \wedge \text{inst}(x, C))) \\ \rightarrow B = C$$

Species of a common genus never share instances.

$$\text{A16. } (\text{genus}(A) \wedge \text{inst}(x, A)) \rightarrow \exists B (\text{lowestspecies}(B) \wedge B \text{ is a } A \wedge \text{inst}(x, B))$$

Every instance also instantiates some lowest species.

The above are non-trivial. They have the following theorems as consequences:

$$\text{T14. } \text{genus}(A) \rightarrow \exists B \exists C (\text{nearestspecies}(B, A) \wedge \text{nearestspecies}(C, A) \wedge B \neq C)$$

Every genus has at least two nearest species.

$$\text{T15. } (\text{genus}(A) \wedge \text{lowestspecies}(B) \wedge \exists x(\text{inst}(x, A) \wedge \text{inst}(x, B))) \rightarrow B \text{ is a } A$$

If an instance of a lowest species instantiates some genus, then the lowest species is subsumed by the genus.

$$\text{T16. } \text{nearestspecies}(A, B) \\ \rightarrow \neg \exists C (\text{nearestspecies}(A, C) \wedge \text{nearestspecies}(C, B))$$

If A is a nearest species to B , then there is no path through the hierarchy from B to A via some third class C .

$$\text{T17. } \text{class}(A) \wedge \text{class}(B) \\ \rightarrow (A = B \vee A \text{ is a } B \vee B \text{ is a } A \vee \neg \exists x(\text{inst}(x, A) \wedge \text{inst}(x, B)))$$

Distinct classes are either such that one subsumes the other or they have no instances in common.

To prove further desirable theorems we would need to add an **additional axiom** to the effect that the universe is finite (in other words that there are only finitely many biological classes, and only finitely many instances of such classes), a thesis which seems intuitively plausible in the domain of biology. We could then infer:

$$\text{T18. } (\text{genus}(A) \wedge \text{genus}(B) \wedge \exists x(\text{inst}(x, A) \wedge \text{inst}(x, B))) \\ \rightarrow \exists C (C \text{ is a } A \wedge C \text{ is a } B)$$

If two genera have a common instance then they have a common subclass.

$$\text{T19. } A \text{ is a } B \wedge A \text{ is a } C \rightarrow (B = C \vee B \text{ is a } C \vee C \text{ is a } B)$$

Classes which share a subclass in common are either identical or one is subordinated to the other.

The system so defined implies that each class hierarchy constitutes a supremum semilattice, or in other words that every collection of classes has a least upper bound with

respect to *is a*. To generate a simple model let P be any finite set. P can be, for example, a finite subset of the natural numbers. Let I be any non-empty proper subset of P and let C be any non-empty subset of $\wp(I)$ (so that C is a collection of subsets of I), with the following properties:

- i) $\emptyset \notin C$
- ii) $\cup C = I$
- iii) if $X, Y \in C$ and $X \cap Y \neq \emptyset$, then $X \subseteq Y$ or $Y \subseteq X$
- iv) if $X \in C$, then there is some $Y \in C$ such that either $X \subset Y$ or $Y \subset X$
- v) if $X, Y \in C$ and $X \subset Y$, then there are $Z_1, \dots, Z_n \in C$ disjoint from X such that $X \cup Z_1 \cup \dots \cup Z_n = Y$.

The particulars in these models are the members of P , universals are the members of C , and instantiation is interpreted as the set-membership relation. Instances, then, are the members of I and all members of C are not merely universals but also classes. A highest genus is a member of C that is not a proper subset of any member of C and a lowest species is a member of C that is not a proper superset of any member of C . Notice that because I is finite, there must be at least one highest genus and more than one (but only finitely many) lowest species.

Aristotelian Definitions

We can now, again following Berg (*op. cit.*) give an account of the Aristotelian theory of definitions. To give a definition, for Aristotle, is to say of something what it is. More precisely, a definition tells us what makes an entity of a given sort an entity of that sort. In a different terminology, an Aristotelian definition is an account of the essence or nature of something. Definitions, for Aristotle, are *real* rather than merely *nominal* definitions: thus they are not the specifications of the meanings of words.

It follows from the above that only what has an essence can be defined, and it is precisely classes, in the terms we have been using above, which satisfy this condition. More precisely, it is *species* which can be defined, via the specification in each case of the relevant nearest genus and differentia. The latter tells us what marks out instances of the species within that genus. Thus *human* is defined as *rational animal*, where *animal* is the genus and *rational* is the differentia. The differentia is also referred to in Aristotelian terms as the ‘specific difference’ or ‘difference that makes a species’.

To specify a class is to provide an answer to a “What is it?” question. When faced with a new kind of biological phenomenon the task of the biologist is to provide the tightest possible answer to the “What is it?” question, which means: to provide the species for the phenomenon, which means also specifying the relevant nearest genus and the relevant specific difference. An Aristotelian definition must then satisfy the condition that an entity satisfies it *if and only if* it instantiates the corresponding species. Specifying a genus alone would be to provide an answer to the “What is it?” question that is not sufficiently tight, since the genus encompasses also other phenomena. Note that to specify the qualities, functions or roles of entities or to say what processes entities engage in is not to provide an answer to the “What is it?” question.

Differentia, too, are universals in the sense of this term presupposed in the above. Differentia are not *instantiated*, but rather *exemplified*, a new primitive notion which we symbolize by means of *exemp*. An Aristotelian definition then has the form:

An *A* is a *B* which exemplifies *S*

where the variables *S*, *T*, ... range over differentia. We then have:

$$\mathit{inst}(x, A) \leftrightarrow \mathit{inst}(x, B) \wedge \mathit{exemp}(x, S)$$

and we can define what it is to be a differentia as follows:

$$\text{D9. } \text{differentia}(S) =_{\text{def}} \exists B \exists C (\text{nearestspecies}(B, C) \wedge \forall x (\mathit{inst}(x, B) \leftrightarrow (\mathit{inst}(x, C) \wedge \mathit{exemp}(x, S))))$$

The genus together with the differentia of a species constitutes the *essence* of the corresponding species.¹⁸

We can then postulate axioms for differentia such as:

$$\text{A17. } \text{differentia}(S) \rightarrow \neg \text{class}(S)$$

and prove theorems for example to the effect that:

$$\text{T20. } \text{differentia}(S) \rightarrow \exists x \mathit{exemp}(x, S)$$

The axioms presented above are motivated by the sorts of classifications we find in the life sciences. However, in extensions of the theory we may consider which amended versions of these axioms would be required to cope with the classification of natural kinds in non-organic domains such as chemistry, meteorology, or geomorphology,¹⁹ and also which axioms, or systems of axioms, would be needed to cope with the classification of *artefacts* of different sorts, including both physical artefacts such as drugs or drug-delivery devices, and non-physical artefacts such as medical procedures or diagnoses. We may extend the framework still further by considering what, if any, would be the analogues of the axiom systems here considered in realms such as temperature, which are marked by continuous variation, or in realms such as types of soil or types of water impurity, which are marked by combinations of factors which vary independently. We may consider what the analogues of these axioms would be for the different sorts of folk classifications carried out by human beings in different cultures and using different natural languages.²⁰ And finally, and most importantly for the realm of biomedical

¹⁸ Not everything which satisfies D9 is a differentia, for Aristotle, who distinguished also what he called ‘*propria*’, which are properties peculiar to all the members of a given species which yet do not belong to the essence of the species – for example the property: *capable of laughing* as possessed by humans. We ignore this issue here for the sake of simplicity.

¹⁹ Barry Smith and David M. Mark, “Do Mountains Exist? Towards an Ontology of Landforms”, *Environment and Planning B (Planning and Design)*, 30(3), 2003, 411–427.

²⁰ See Douglas L. Medin and Scott Atran (eds.), *Folkbiology*, Cambridge, MA: MIT Press, 1999.

informatics, we may consider how to manipulate simultaneously a multiplicity of different classifications, prepared in different disciplinary contexts or for different purposes, of the same domain of phenomena in reality.²¹

The Foundational Model of Anatomy

Among all existing biomedical ontologies it is the Foundational Model of Anatomy, developed at the University of Washington, Seattle as part of the Digital Anatomist Project, which comes closest to meeting the standards of formal rigor taken for granted among philosophical ontologists. The Foundational Model of Anatomy (hereafter: FMA) is a symbolic representation of the structural organization of the human body from the macromolecular to the macroscopic levels.²² It has the goal of providing a robust and consistent scheme for classifying anatomical entities on the basis of explicit definitions of a sort which can serve as a reference ontology in biomedical informatics.

Most significant, from our present point of view, is the fact that the FMA has adopted an Aristotelian regime of definitions.

Thus definitions in FMA look like this:

Cell *is a* **anatomical structure** that *consists of* **cytoplasm** *surrounded by a* **plasma membrane** with or without a **cell nucleus**

Plasma membrane *is a* **cell part** that *surrounds* the **cytoplasm**

where terms picked out in bold are nodes within the FMA classification and italicized terms signify the formal relations – including *is a* – which are defined between these nodes.

As the FMA points out, ontologies ‘differ from dictionaries in both their nature and purpose.’²³ Dictionaries are prepared for human beings; their merely nominal definitions can employ the unregimented resources of natural language, can tolerate circularities and all manner of idiosyncrasy. In ontologies, however, definitions must be regimented in such a way that each reflects the position in the hierarchy to which the definiendum belongs:

The role of definitions in an ontology is ... to specify such defining attributes in a consistent manner, thus assuring their transitive inheritance through a type hierarchy. Consistency in definitions and, therefore, in the classification, requires that a unifying viewpoint (i.e., context) be also specified for concept representation.

²¹ See Thomas Bittner and Barry Smith, “A Theory of Granular Partitions”, *Foundations of Geographic Information Science*, Matthew Duckham, Michael F. Goodchild and Michael F. Worboys, eds., London: Taylor & Francis, 2003, 117–151.

²² See Cornelius Rosse and José L. V. Mejino Jr., “A Reference Ontology for Bioinformatics: The Foundational Model of Anatomy”, *Journal of Biomedical Informatics*, forthcoming.

²³ J. Michael, José L. V. Mejino Jr., and Cornelius Rosse “The Role of Definitions in Biomedical Concept Representation”, *Proceedings of the American Medical Informatics Association Symposium*, 2001, 463–467.

This context should hold true for the entire ontology.²⁴ Provided such requirements are satisfied, the position of a concept will enrich its own definition by the definition of all of its parents within the hierarchy. Thus, unlike in a dictionary, a definition of a concept within an ontology is incomplete without that of all of its parents.²⁵

This means additionally that in order to ensure transitive inheritance of essential characteristics, all intermediate classes should be defined even if they have not have been explicitly identified in the scientific literature.²⁶ It means also that, already on the basis of its rules for the formulation of definitions, the FMA rules out multiple inheritance. And it means, finally, that the FMA, with its system of definitions, can exploit all the benefits – in terms of reliable curation, efficient error checking and information retrieval, and ease of alignment with neighboring ontologies – of *logical compositionality*.

GO Again

GO, too, like other, similar biomedical ontologies, provides not only controlled vocabularies with hierarchical structures but also definitions of its terms. Indeed part of the goal of GO, and of similar projects, is to provide a source of ‘strict definitions’ that can be communicated across people and applications. When we examine GO’s actual practice, however, we find that its definitions are affected by a number of characteristic problems which, while perhaps not affecting their usability by human biologists, will raise severe obstacles at the point where the sort of formal rigor needed by computer applications (or by a formally rigorous biology of the future) is an issue. Consider again our two initial examples:

GO:0003673: cell fate commitment

Definition: The commitment of cells to specific cell fates and their capacity to differentiate into particular kinds of cells.

GO:0045168: cell-cell signaling involved in cell fate commitment

Definition: Signaling between cells that results in the commitment of a cell to a certain fate. This is often done by secretion of proteins by one cell which affects the neighboring cells and causes them to adopt a certain fate.

In both of these definitions we recognize the characteristic problem of circularity. The coarse logic of the definition of cell fate commitment is as follows:

x is a cell fate commitment $=_{\text{def}}$ x is a cell fate commitment and p ,

where p is, logically speaking, a second, extraneous condition. Further problems arise in virtue of the fact that, as a result of its use of unregimented natural language and of its

²⁴ The context for the FMA is: anatomical structure; this means that all the definitions in the FMA hierarchy are formulated exclusively in structural terms, which means: without appeal to normal and abnormal functions performed by the anatomical entities distinguished.

²⁵ *Op. cit.*

²⁶ *Op. cit.*

lack of concern for issues of logical compositionality, substitution of GO definiens for the GO terms appearing within other GO terms and definitions can be achieved, at best, only with human intervention. Thus consider:

GO:0030154: Cell differentiation

Definition: The process whereby relatively unspecialized cells, e. g. embryonic or regenerative cells, acquire specialized structural and/or functional features that characterize the cells, tissues, or organs of the mature organism or some other relatively stable phase of the organism's life history.

GO:0007514: Garland cell differentiation

Definition: Development of garland cells, a small group of nephrocytes which take up waste materials from the hemolymph by endocytosis.

In this way a number of valuable methods of inference, extrapolation of new terms, and error-checking are foreclosed.

Conclusion

The treatment of *is a* relations in biomedical ontologies has been thus far highly problematic. In some cases, indeed, the two relations are not clearly distinguished at all, leading to what Guarino calls '*is a* overloading'.²⁷ The FMA defines an ontology as a 'true inheritance hierarchy',²⁸ thereby drawing attention to the fact that one central reason for adopting the method of ontologies in supporting reasoning across large bodies of data is precisely the fact that this method allows the exploitation of the inheritance of properties along paths of *is a* relations.

When challenged with such problems, the members of the GO and associated communities standardly insist that their concerns are those of practicing biologists, and that they are thus not concerned with the sorts of scrupulousness that are important in logic. To repeat, however, if GO's adherents propose that GO should serve as a reference-platform for computer-assisted navigation between biomedical databases, then the failure to achieve consistency with standard logical principles will place considerable obstacles in the way of its efforts to achieve this end.²⁹

²⁷ Guarino, N. "Some Ontological Principles for Designing Upper Level Lexical Resources", in *Proceedings of the First International Conference on Language Resources and Evaluation*, Granada, 1998, 527–534.

²⁸ Michael, Mejino, and Rosse, *op. cit.*

²⁹ Acknowledgements: This paper was written under the auspices of the Wolfgang Paul Program of the Alexander von Humboldt Foundation. Thanks are due also to Maureen Donnelly, Kai Hauser, Ingvar Johansson, Jacob Köhler, David Mark, Carsten Pontow, Cornelius Rosse, Steffen Schulze-Kremer and Jonathan Simon.