

# The Relationship of Philosophy to Science in the Information Age: a proposal for understanding Information Ontology

Andrew D. Spear

## Introduction

There has been a growing consensus in certain quarters over the last twenty years that something called ‘ontology’ is going to play an increasingly important role in the future of scientific research and in representing, consolidating and making widely available the massive and ever-growing amounts of scientific information, especially in the bio-medical sciences.<sup>1</sup> The word ‘ontology’ here has a double meaning. In philosophy ‘ontology’ has the meaning, roughly, of being the study of being, especially of the kinds or categories of being that there are and of the very general or formal relationships that obtain amongst these.<sup>2</sup> In computer science, ‘ontology’ has come to mean, roughly, a very general set of interconnected terms that are to be used to organize information from a given domain. These two senses of the term are not unrelated and it is the needs of twenty-first century scientific information management that are now bringing these two understandings of ontology together.

The big idea can be outlined as follows. There is now more information available in the bio-medical sciences than could possibly be learned, memorized and reasoned about by any single human individual, or even by a large research group composed of many such individuals; and yet there is a great need on the part of working scientists, as well as working doctors, to be able to efficiently access and use scientific information that is outside of their areas of specialization, or that has been recently produced as the result of new research. The obvious solution to this problem is to make use of the super-human memory and reasoning capacities of computers. Ideally, if a doctor wants to know the latest scientific information on a kind of cancer she is treating she should be able to look it up in a well organized and regularly updated “cancer database” of some sort. The major obstacle to this has proven to be that computers are not, *prima facie*, able to navigate differences in native language, terminology-use and semantic content of the scientific information that they are keeping track of in the way that human beings

---

<sup>1</sup> See for example Barry Smith, "Ontology," in *Blackwell Guide to the Philosophy of Computing and Information*, ed. L. Floridi (Oxford: Blackwell, 2003).

<sup>2</sup> For a general discussion of Metaphysics and ‘Ontology’ in this philosophical sense, see the opening chapters of E. J. Lowe, *The Four-Category Ontology: A Metaphysical Foundation for Natural Science* (Oxford: Oxford University Press, 2006), ———, *The Possibility of Metaphysics: Substance, Identity and Time* (Oxford: Oxford University Press, 1998).

are; nor are computers able to represent the internal structure and organization of scientific theories in the ways that human scientists and textbooks do unless all features of this internal structure and organization are explicitly included in the computers' programs. But representing the internal structure and organization of scientific theories in a terminologically consistent and systematic way just is stating the kinds of objects and relations amongst them that such theories are committed to, and providing a general set of interconnected terms for the purpose of organizing information just is the task of ontology (or at least some part of it) in both the philosophical and the computer science senses of the term. The upshot of all of this is that the advent of the computer age in conjunction with the prodigious and unceasing output of the sciences have combined to make philosophical ontology extremely relevant to the tasks of representing, organizing and retrieving scientific information in the twenty first century.

This conception of ontology, which might be called "information ontology", raises both traditional philosophical issues about the categories of being, the nature of representation and the relationship of philosophy to science, and also issues that have arisen more recently, such as the possibility and nature of artificial intelligence and the absolute limits on expressivity of a formal theory relative to its computability in a machine or computer context. My goal here is to provide a very general definition of 'information ontology', one that will give rise to and itself be formulated in terms of a framework within which some of these more specific issues can be addressed. An "ontology", in the sense of the kind of discipline that will play a key role in dealing with issues of information management in the twenty first century has recently been defined as,

"a representational artifact whose representational units are intended to designate universals in reality and the relations between them."<sup>3</sup>

Beginning with the insight expressed in the first part of this definition, I will go a step further and suggest that 'ontologies' in the information science sense be understood as "computational systems" in the sense of David Marr, and that they are thus susceptible to analysis in terms of what Marr and later Michael R. W. Dawson have called the "tri-level hypothesis".<sup>4</sup> Having established this basic framework, I will address the relationship of information ontologies to the

---

<sup>3</sup> Barry Smith et al., "Towards a Reference Terminology for Ontology Research and Development in the Biomedical Domain." (paper presented at the KR-Med, 2006).

<sup>4</sup> See David Marr, *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information* (New York: W. H. Freeman and Company, 1982), M. R. W. Dawson, *Understanding Cognitive Science* (Oxford: Blackwell Publishers, 1988).

question of artificial intelligence, propose an extremely general definition of what an information ontology is, and then discuss the relationship of this definition (against the background of the tri-level hypothesis) first to the traditional philosophical issues of universals and of the nature of representation, and then, much more briefly, to issues of expressivity vs. computability, and of the constraints placed on information ontologies by the way in which they are implemented in particular computing devices. The result, I hope, will be a framework that might serve as a basis for discussion in evaluating the relevance of philosophy generally to issues of information science, as well as serve as a basis for interdisciplinary communication, and for the comparison and evaluation of different ontologies built on different philosophical assumptions and computational methods.

### Information Ontologies as Representational Artifacts as Computational Systems:

#### The tri-level hypothesis

Ontologies are representational artifacts and ‘ontologies’ in the information science sense of the word are “formalized” or computationally tractable representational artifacts.<sup>5</sup> They are thus analogous to what Michael R. W. Dawson, following David Marr and others have called “computational systems”. However, not all formalized representational artifacts are ontologies. Relying on the idea of ontologies as computational systems, we will here distinguish two major dimensions along which a representational artifact can be considered an ontology: the semantic interpretation of the system on the one hand, and the formal or logical properties of the semantic interpretation on the other.

A computational system is an information processing system that consists of representational units or symbols, a set of rules for the manipulation of these symbols and a semantic account or intended interpretation that explains what the symbols mean or refer to, and what it is that is happening when these symbols are manipulated according to the rules of the system. Developing ideas from David Marr, Michael R. W. Dawson has proposed what he calls the “tri-level hypothesis”, the thesis that computational systems of this sort involve three distinct levels of analysis and explanation. These are the semantic level, the algorithmic or computational level and the implementational level.

#### *The Semantic Level*

---

<sup>5</sup> On some of the basic terminology being used here, see once again Smith et al., "Towards a Reference Terminology for Ontology Research and Development in the Biomedical Domain."

The semantic level of analysis is the level at which the intended interpretation of the representational units and rules for their manipulation is determined. It is the level at which a computational system is understood as representing certain kinds of things (such as concepts, diseases, numbers, patients in an emergency room, classes or universals in reality) and of manipulating these representations in certain specific ways (tracing relationships amongst concepts, providing stored definitions of diseases, adding numbers together to provide a sum, storing and retrieving information about patient history and treatment, navigating is\_a hierarchies amongst classes or universals). The semantic level is crucial for the use and interpretation of computational systems by human beings and it is thus important that computational systems intended for use by human beings have clear, understandable and accurate semantic interpretations assigned to them.

#### *The Algorithmic Level*

The algorithmic level is the level of programming, of software construction and analysis. The question at the algorithmic level is that of what kinds of computer tractable languages and rules to use in order to properly capture the intended interpretation of the information system, and in order to ensure effective and accurate inferences. It is important to recognize that many different programs or algorithms can be understood as representing the same information and performing the same semantic tasks, and that the same algorithm can be interpreted as representing different information and performing different semantic tasks. There is thus no simple isomorphism between the algorithmic and the semantic levels of an information system, though the two do place certain limitations and constraints on one another.

#### *The Implementational Level*

The implementational level is the level of physical instantiation or hardware. Implementation has to do with the physical device and processes that are used to “run” or instantiate the algorithms of a computational system. While most computers run off of silicon chips, many cognitive scientists believe that there is a strong resemblance between the working of computers and the functioning of the human mind, and that therefore neurons are another candidate for the physical implementation of certain kinds of computational systems. What is most crucial in understanding a computational system is to understand its semantic and algorithmic components. However, the implementational level also places some constraints on what information can be represented, and in what way, insofar as it is in part properties of the

physical instantiation of a computational system that will determine its processing speed and “memory capacity”.

Important for our purposes here is to recognize that ontologies in the information science sense of the term are a kind of computational system, and that designing an information ontology can thus be understood in terms of the tri-level hypotheses as i) specifying the things to be represented and the tasks to be performed with the represented information at the semantic level, ii) determining an appropriate programming context or environment for these representations and tasks at the algorithmic level, and then iii) implementing this programming in a physical device (computer) that possesses sufficient processing speed, memory, etc. to actually perform the representation and information processing tasks that are intended at the semantic level. It is crucial at the outset to point out that the primary determinant as to whether or not a computational system is an ontology will be the features of its intended semantic interpretation, and only very secondarily the features of its programming and implementation.

#### Do information ontologies think?

An important initial issue that needs to be addressed in connection with understanding information ontologies as a kind of computational system has to do with what it means for the representational units and rules of a system to “have an interpretation” or semantic account. In what sense are the representations in a computational system “about” anything, and in what sense can the system be understood as “doing something”? There are two possible answers to this question.

The first, popular with philosophers and other researchers interested in the possibility of artificial intelligence, is that computational systems are “about” reality and “do things” or “perform tasks” in the same sense that we human beings do. In other words, computational systems have, or can have, minds and thoughts just like we do. Thus, these systems have contentful thoughts and beliefs (representational units) that are *about* or directed at things in the world. To understand the semantics of computational systems in this way is to ascribe to them what John Searle has called intrinsic or original intentionality, and to view them as minds in the literal sense of the word.<sup>6</sup>

---

<sup>6</sup> John R. Searle, *Intentionality: An Essay in the Philosophy of Mind* (Cambridge: Cambridge University Press, 1983).

The second and alternative answer is to say that computational systems have semantic accounts only in a secondary or derivative sense, based on the interpretations that we (thinking human subjects) impose on the representational units they contain and the manipulations of these units that they perform. This is to ascribe to computational systems what John Searle has called “derivative intentionality”, and to maintain that these systems are only “about” something or “processing information” insofar as some other intentional subject has assigned meanings to the machine’s representational units and manipulations thereof.<sup>7</sup>

The difference between these two answers amounts to answering the question “do computers think like human beings do?” with a “yes” in the first case, and with a “no” in the second. While we do not wish to take a definitive stance on the question of whether or not it is in principle possible or impossible for computers to think and/or be conscious in the same way that human subjects are, it is crucial for understanding contemporary work on information ontologies that the semantic interpretation of the representational units in these systems involves *derived*, not *original*, intentionality. The Gene Ontology, the Foundational Model of Anatomy and other such information ontologies do not have minds, do not have thoughts and beliefs, and do not contain concepts *in the same way* that human subjects do.<sup>8</sup> The Gene Ontology is not a biologist, nor is the Foundational Model of Anatomy an anatomist. These systems cannot simply communicate with each other in natural or some other language in order to align the information that they contain, they cannot systematically analyze the information that they contain for definitional adequacy and logical consistency, they cannot check the information that they contain against their perceptions of reality, and they cannot have a conversation with a human user in which they explain, disambiguate and answer questions about the intended interpretation of their information content, *unless* they have been explicitly and systematically designed by actual intentional human subjects to do so. It is precisely this fact about all currently existing information systems that leads to problems of ambiguity, poor definitions, use-mention confusion, information integration etc. that have been discussed extensively in recent literature.<sup>9</sup> Thus any kind of

---

<sup>7</sup> For more on derived and original intentionality, see Ibid.

<sup>8</sup> The Gene Ontology can be found at <http://www.geneontology.org/>, while the Foundational Model of Anatomy is on-line at <http://sig.biostr.washington.edu/projects/fm/AboutFM.html>.

<sup>9</sup> Jacob Kohler et al., "Quality Control for Terms and Definitions in Ontologies and Taxonomies," *BMC Bioinformatics*, no. 7 (2006), <http://www.biomedcentral.com/1471-2105/7/212>, Barry Smith, Jacob Kohler, and Anand Kumar, "On the Application of Formal Principles to Life Science Data: A Case Study in the Gene Ontology," in *Database Integration in the Life Sciences (Dils 2004)*, ed. E. Rahm (Berlin: Springer, 2004).

computational system requires a specific and rigorous statement of its intended interpretation and use, and this goes not least for information ontologies themselves.

#### A general definition of ‘information ontology’

We have already defined, following Smith et al., an information ontology as a representational artifact whose representational units are intended to designate universals in reality and the relations between them.<sup>10</sup> This definition implicitly presupposes a realist semantic interpretation of the representational units in an ontology. However, for purposes of exposition and for the sake of establishing comparisons and contrasts, we provide the following more general definition of an information ontology:

An information ontology is a *representational artifact* the intended interpretation of which is *what is general* in reality, and the *relationships amongst what is general*, and which includes a *maximal degree of precision* regarding the features of what is general and of the relationships between these features.

Based on this definition, whether or not something is an ontology can be determined along two basic dimensions.

The first is the dimension of *what* the ontology represents. The schematic definition we have given here rules out representation of that which is particular, individual or an “instance” in reality as an intended interpretation for an information ontology, but leaves open the question of whether what is general consists of universals, classes, concepts or arbitrarily constructed equivalence classes that have been attached to the conventional uses of words in natural language; issues which we will discuss in the next section.

The second dimension along which something can be evaluated as to whether or not it is an ontology is in the explicitness and detail with which it represents what it represents. The detail, precision and rigor of an ontology can be measured in terms of the cogency of its definitions, and the logical features of the structures (such as taxonomies) that it represents, however it is important to distinguish between features of the ontology as a representation (such as its logical features) and what those features represent on the side of reality. Just because the axioms and/or logical system defining an ontology are sound, consistent or incomplete does not mean that the things that they represent also have these features (is there such a thing as a “sound” or

---

<sup>10</sup> Smith et al., "Towards a Reference Terminology for Ontology Research and Development in the Biomedical Domain."

“complete” group of objects or structures in reality?). It is thus the logical features of the semantically relevant or interpreted information in a representation that will play a primary role in determining whether or not that representation is an ontology or some other kind of computational system.

### Semantic Interpretation of a Representational Artifact

What do (or should) the representational units and relationships amongst them in an ontology refer to? We have already ruled out particulars or individual entities in the world, by saying that ontologies represent “what is general” in reality. However, as both the history of philosophy and recent debates in information science show, how to interpret this phrase “what is general” is a matter of much controversy. It is also a matter of great controversy whether the representational units in an ontology represent what is general in reality directly, or whether they represent it indirectly, in virtue of directly representing concepts or conceptual schemes belonging to some person or group of people. Here we will briefly lay out the options and sub-options associated with each of these views of what it is that ontologies represent.

#### *The representational units in an ontology represent what is general in Reality*

The question of what it is that is general in reality is roughly the question of what it is that makes scientific information, generalizations and law-like statements true. What is it that general terms such as ‘eukaryotic cell’, ‘H<sub>2</sub>O’ and ‘cancer’ refer to, and what is it that all things that these words apply to have in common that make them things of this kind?

#### *Metaphysical Realism*

The metaphysical realist answer to this question is that what is general in reality are universals, mind-independent multiply instantiated or realized entities such as the species “cat”, the color “red” and the element “oxygen” that are the referents of the general terms ‘cat’, ‘red’, ‘oxygen’ etc. that we use, and that serve to ground the regularity and order that is exhibited by these kinds in nature. Historically, realism comes in two versions: Platonism and Aristotelianism. According to Platonism, universals are dependent neither on the minds of individuals nor on the existence of their instances, but can exist and be what they are independently of whether or not any other objects exist at all. Aristotelianism or ‘immanent realism’ generally holds that there are mind-independent universals, but that they exist only as instantiated in their instances. There are nuances to and multiple variations of both Platonism and Aristotelianism, however what has been said so far should already serve to make the basic

proposal of the realist about universals clear. The basic idea behind metaphysical realism is that there are facts about reality that ground or make true certain generalizations and judgments of both science and common sense, and that render other such generalizations and judgments false. According to metaphysical realism, reality has *bona fide* general or abstract joints, and it is the task of knowing subjects, especially those engaged in the pursuit of natural science, to discover what these are.

### *Nominalism*

The major alternative to realism about universals is nominalism. Nominalism is effectively the claim that universals do not exist, and that there are no special entities responsible for the order and regularity that nature displays, or for the fact that things belonging to a kind exhibit similarity with respect to certain of their features or properties. The most prominent historical nominalist may well be William of Ockham, whose writings gave rise to the famous slogan known as “Ockham’s Razor”, which states “don’t multiply entities beyond necessity”.<sup>11</sup> Nominalists characteristically argue that the general terms that we use pick out classes or sets of entities in reality that we have grouped together based on perceived resemblances and then attached to either a concept in the mind, to the conventional use of a word, or both. According to nominalism, there may be some *bona fide* joints or divisions in reality, but for the most part our judgments and generalizations involve us in imposing order and regularity on a reality that does not possess such structure independently of us, of our language use and of our conceptualizations.

### *Metaphysical Conceptualism (?)*

Philosophers often speak of a third great historical alternative to realism and nominalism: “conceptualism”. Conceptualism can be defined, roughly, as the view that universals do not exist in the mind-independent world, but that they do exist as concepts in the minds of human subjects. Conceptualism is an uncomfortable cousin to realism and nominalism however; as it seems that it is either a version of realism, or a version of nominalism, or else an answer to the question of what is general in reality that is highly implausible.

Most Aristotelian realists are also conceptualists in the following sense: Aristotelian realists construe universals as existing only in the particulars that are their instances, however,

---

<sup>11</sup> Spade, Paul Vincent, "William of Ockham", *The Stanford Encyclopedia of Philosophy (Fall 2006 Edition)*, Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2006/entries/ockham/>>.

they also maintain that it is possible for the human mind to gain knowledge of universals. This happens when a human subject, after examining many instances of the universal in question (examining many prokaryotic cells, for example) comes to form a general idea or concept in their mind that corresponds to the universal that all of those individuals instantiate. This is a perfectly plausible picture of how knowledge of universals is gained on a realist account (indeed, something like it seems to have been Aristotle's view), however it is not clear why such a position should be called "conceptualism", as opposed to "Aristotelian realism" full stop, since the role of concepts on such a view is only to explain our knowledge of what is general in reality, not to account for what is general in reality itself.

On the other hand, if the conceptualist claim that universals do not exist in reality independently of the minds of thinking subjects is to be taken seriously, then it is difficult to see how this position is different from most versions of nominalism, insofar as nominalists say the same thing, and usually add an account of what it is that general terms do refer to that involves concepts, sets and/or language use of some sort on the part of cognitive subjects.

Finally, conceptualism might be construed as a genuine alternative to realism and nominalism in the following way: the conceptualist thesis could be that what accounts for the generality, structure and order in nature literally *is* the existence of certain kinds of concepts. Thus, in the absence of any minds with concepts, there would be nothing that all cats have in common or instantiate that makes them all feline, but as soon as some mind with the concept "cat" exists, then suddenly the facts of reality themselves change and the things that were once simply a random class suddenly have features that are ontologically similar in virtue of instantiating the mental universal or concept "cat". This view is independently strange, but still seems close enough to nominalism as to be hardly worth treating as a legitimate third option. Thus, with regard to the question of what it is that accounts for what is general in reality, there are roughly two answers: "something" and "nothing" (or at least nothing over and above particular objects viewed in a certain way), and these correspond rather exactly to metaphysical realism and nominalism respectively.

*The representational units in an ontology represent what is general indirectly by representing  
Concepts*

One of the most popular answers to the question "what do the representational units in an ontology represent?" in contemporary information science is that these units represent

“concepts”.<sup>12</sup> The intuition behind this claim seems to be that knowledge and information consist of concepts, and that therefore representing the general knowledge or information that we possess requires representing the concepts that constitute or comprise this knowledge. In embracing these kinds of positions, contemporary information ontologists implicitly reject what might be called *perceptual* or *epistemological realism*, in favor of one of two other positions that have been traditionally called *representationalism* and *idealism* respectively.<sup>13</sup> Whereas *realism* in this sense states that our perceptions and our thoughts are directly about entities of various sorts (particulars and universals in reality), *representationalism* maintains that our perceptions and thoughts are directly about certain kinds of mental objects (representations, appearances, concepts, etc...) and therefore only indirectly about things in reality, while *idealism* maintains, roughly, that our perceptions and thoughts aren't about reality at all, but are entirely about mental objects such as perceptions, appearances, concepts and ideas. Thus, just as there is a debate regarding the proper understanding of *what it is in reality* that general terms refer to, so also there is a debate about *the way in which* the mind, or a representational device intended to encode meaningful information in a derivate way, succeeds in referring to the things that its thoughts or representations are about.

Now, it is important to note at the beginning that the existence of concepts is consistent with all of the metaphysical positions about universals that were outlined in the previous section. For the realist, individuals will have concepts or general ideas, and some of these concepts will be accurate in virtue of representing or capturing universals on the side of reality, while others will be inaccurate or false in virtue of failing to do so. On many nominalist accounts it is precisely concepts of one sort or another that are the meaning or significance of the general terms that we use, while for a “conceptualist” about universals concepts would be, to say the least, a crucial component of his understanding of reality. Thus the question of whether or not

---

<sup>12</sup> For more on “concepts” and problems with the use of conceptual terminology in information ontology, see Barry Smith, "From Concepts to Clinical Reality: An Essay on the Benchmarking of Biomedical Terminologies," *Journal of Biomedical Informatics*, ———, "Beyond Concepts: Ontology as Reality Representation," in *Proceedings of Fois 2004*, ed. Achille Varzi and Laune Vieu (2004), Gunnar O. Klein and Barry Smith, "Concepts Systems and Ontologies: Recommendations Based on Discussions between Realist Philosophers and Iso/Cen Experts Concerning the Standards Addressing "Concepts" And Related Terms," (2005), <http://ontology.buffalo.edu/concepts/ConceptsandOntologies.pdf>, Ingvar Johansson, "Bioinformatics and Biological Reality," *Journal of Biomedical Informatics* 39 (2006), Barry Smith and Werner Ceusters, "Wüsteria" (paper presented at the The XIXth International Congress of the European Federation of Medical Informatics, 2005).

<sup>13</sup> For an overview of these different positions, see David Kelley, *The Evidence of the Senses : A Realist Theory of Perception* (Baton Rouge: Louisiana State University Press, 1986).

representational units in ontologies represent things in reality or concepts of those things is not primarily a response to the philosophical debate about what it is that accounts for the generality and order that reality seems to exhibit, rather it is a response to the question of *how* the mind and its concepts relate to reality, whether in a realist, representationalist or idealist sort of way.

However, even this does not entirely clarify what it means to say, “The representational units in an ontology are or represent concepts.” On consideration, there seem to be four options as to what this might mean.

The first possibility is to read the statement literally: ‘representational units in ontologies are concepts’ just means that the representational units comprising the information contained in an ontology *are identical with* concepts. This interpretation seems clearly wrong, first because it is unclear exactly what it adds to the project of information ontology if true, but second because it would seem to involve the information ontologist in the claim that he or she has definitively solved a central question of AI as to whether or not computers can think and possess concepts in the same way that humans do in the affirmative. This would be a remarkable accomplishment, but not one that the disciplines of information science or information ontology in their current states are likely to be eligible to lay claim to.

A second way of interpreting the claim that the representational units in an ontology are concepts would be to read it as the claim that “representational units in ontologies *represent* concepts”. This seems closer to the intended meaning and is consistent with some versions of what I have labeled *representationalism*, but it still raises some questions. First, if concepts are representations, then this understanding of information ontologies construes ontologies themselves as representations of representations, and this seems *prima facie* redundant. Second, it isn’t clear exactly what it would mean to represent a concept. Philosophers, psychologists and neuroscientists have been working and still are working quite laboriously in order to gain knowledge of whether or not concepts exist, what they are, how they are acquired and how they function in human thought and action. However, if information scientists are already representing concepts in their ontologies, then they presumably understand what concepts are, indeed, they understand what concepts are well enough to provide characterizations, definitions and perhaps even identity conditions for them. If this is the case, then once again it would seem that information science has managed to make a substantial contribution to our understanding of

the mind without even realizing it or publicly announcing it.<sup>14</sup> Thirdly, and related to the foregoing point, if the representational units in an ontology are genuinely to be interpreted as representing concepts, then the unit in a biological ontology that is intended to represent the kind “cat” is actually intended to represent the concept of “cat”. This means that a full definition of this representational unit should include a definition or statement of identity conditions both for cats and for what it is for something to be the concept of a cat. This would seem to suggest that all information ontologies that claim to be *representations of concepts*, but fail to provide definitions, not only of “cancer” and “eukaryotic cell”, but also for “concept of cancer” and “concept of eukaryotic cell” are fundamentally incomplete relative to their intended semantic interpretations.

A third option for the meaning of the claim that the representational units in an ontology are concepts would be that “the representational units in ontologies represent *what concepts are about in reality in virtue of* representing concepts. This would involve embracing a specifically *representationalist* view of perception and thought, and then applying it to understanding how ontologies work as well. Now, this interpretation seems quite plausible, but it raises the immediate question: if the goal of ontologies is to represent things in reality (what is general), then why go about this indirectly by representing concepts, rather than directly by representing what those concepts are about? It seems that the motivation for this position is probably a sort of epistemological humility on the part of its advocates, and to this extent it is legitimate, even admirable. The intuition seems to run as follows: though we are quite confident, based on the best available evidence and on continuing scientific research, that our theories are true or very close to true of reality itself, and that the laws, relationships, kinds and objects that they posit actually exist, we nevertheless know that many people in the history of science who were confident in a similar way about their own theories later had their theories overthrown, and we ourselves are continuously gaining new scientific information about reality, any bit of which might very well overthrow some portion of one of our currently favored scientific theories, therefore we should not claim that our ontologies represent reality (since they could be wrong),

---

<sup>14</sup> This raises the point that there is, of course, room in a realist view of ontologies for the representation of concepts: specifically in the domain of psychology, where the domain of the ontology will be that of the mental, and the real entities in reality that are to be represented by the representational units in the ontology will be mental types such as belief, emotion, memory and concept.

but rather that they represent our conceptualizations of reality (since even if these conceptualizations are wrong, it will still be true that our ontologies represent *them*).

Once again, this line of thinking certainly has some legitimacy. However, realism (the position, roughly, that our scientific discourse refers, by and large, to things that exist in mind-independent reality) is perfectly compatible with fallibilism (the recognition that our current “best” scientific theories may well have to be revised in light of new evidence and future research). So, presuming that people who maintain this kind of “concept” view of ontologies are in fact committed to capturing truths about reality, the question that immediately presents itself is: why not eliminate the conceptual middle man? We have concepts, and we believe that our concepts refer to things in reality, but also that some of these beliefs may be false. However, computational systems (lacking as all current ones seem to, intrinsic intentional states) do not have concepts, they have only representational units with a (by us) intended interpretation. If our goal is to represent in an ontology what our concepts themselves represent, then we should not seek to represent our concepts, but rather the things that we believe that these concepts are about. If we do this, then it may still turn out that some of our own beliefs fail to refer to things in reality, but then the ontologies that we have built will fail for exactly the same reasons, and can be revised by us in accordance with the revisions that our beliefs (and concepts) about reality have themselves undergone in light of the new scientific evidence.<sup>15</sup>

Alternatively, suppose that an ontology is a representation of our conceptualization of things on the side of reality. Then, even if our conceptualization changes in the future, the ontology remains an accurate, true or faithful representation of our earlier conceptualization, if it ever was, even if we now possess a new and different conceptualization. Ontologies the intended interpretation of which is that of what is general in reality can be evaluated for adequacy relative to the facts of reality, but ontologies the intended interpretation of which is a “conceptualization” can be evaluated for accuracy only relative to that conceptualization, and thus can remain accurate or “true” to the facts of the conceptualization, even if it has been learned that the conceptualization itself is false of reality. Since the goal of scientific ontologies

---

<sup>15</sup> For a discussion of evaluating ontology errors along these lines, see Werner Ceusters and Barry Smith. "A Realism-Based Approach to the Evolution of Biomedical Ontologies." (Place Published, 2006), <http://ontology.buffalo.edu/bfo/Versioning.pdf>. For a highly skeptical treatment of the issue of realism and theory evaluation, see T. S. Kuhn, "Objectivity, Value-Judgment and Theory Choice," in *Scientific Knowledge: Basic Issues in the Philosophy of Science*, ed. J. A. Kourany (Belmont: Wadsworth Publishing Co, 1986). For a realist conception of ontology and its relationship to science, see especially the introduction and chapter 8 of Lowe, *The Possibility of Metaphysics: Substance, Identity and Time*.

is the representation of facts about reality, the question would be why bother with representing facts about conceptualizations at all?

A fourth, final and very extreme interpretation of the thesis that ontologies represent concepts would be to say that representational units in ontologies represent concepts and what they are about,<sup>16</sup> but what concepts are about tells us nothing of significance about reality. To adopt such an interpretation is to adopt an *idealist* or neo-Kantian view according to which the fact that we can only know reality via our thinking and conceptualizing faculties entails that we cannot think about or conceptualize reality *as it is*. Or stronger still, the thesis that what there is in reality depends in some sense on our concepts and our ways of thinking and knowing. There is a long and detailed history of debate in philosophy about whether we can have knowledge of an “external world”, and about whether the universe is fundamentally mind or matter, and whether or not what exists is in some sense dependent on what we think or believe exists. We will not review this history and these debates here. We will simply make a few selected points in passing.

The first is that embracing a strong idealist conception of knowledge and reality such as this one is strongly at odds with the interpretation that most working scientists and even ordinary people have of the results of the physical, biological and medical sciences. It would be odd, to say the least, if computer scientists and information ontologists were to take it upon themselves in the course of assisting natural scientists with the representation of their domain knowledge to both fundamentally reconstrue the intended interpretation of that knowledge, and to advance radical and counter-intuitive philosophical theses about the nature of knowledge and reality as definitively true.<sup>17</sup>

Second, the very idea that is being proposed here is not clear. What it means to say that information ontologies represent concepts and the contents of concepts that have no connection

---

<sup>16</sup> This may well be the view of Gruber, who states that “An ontology is a specialization of a conceptualization... a description (like a formal specification of a program) of the concepts and relationships that can exist for an agent or a community of agents. This definition is consistent with the usage of ontology as set-of-concept-definitions, but more general. And it is certainly a different sense of the word than its use in philosophy.” See Gruber, T. R. n.d. “What is an Ontology?” (<http://www-ksl.stanford.edu/kst/what-is-an-ontology.html>). On this same topic, see also Klein and Smith, “Concepts Systems and Ontologies: Recommendations Based on Discussions between Realist Philosophers and Iso/Cen Experts Concerning the Standards Addressing “Concepts” And Related Terms.”, Smith, “Beyond Concepts: Ontology as Reality Representation.”

<sup>17</sup> Admittedly, the evidence for this claim about the way in which ordinary people and practicing scientists generally interpret the results of science is anecdotal. It would be interesting to see the results of a more systematic survey of the opinions of practicing scientists regarding this matter.

to reality itself is unclear, and any advocate of such a view owes more than some passing remarks in its clarification and defense.

Third and finally, if the Kantian argument that we cannot know reality in itself, but only as our cognitive and conceptual faculties allow us to know it, is being appealed to in order to support this view of information ontology (and it very often is, even if tacitly), then the advocate of this position must answer the following question: in what way do we know about our cognitive and conceptual faculties such that *we can model or represent them*? The answer must be either that we are able to think about them directly, or (to be consistent with the line of reasoning that gets us to this position in the first place) that we can only have knowledge of our cognitive and conceptual faculties indirectly via our conceptualizations of them, and therefore cannot know them in themselves either.<sup>18</sup> But if this is true, then as information ontologists we must seek to model conceptualizations of conceptualizations of reality (since we can no more know our own conceptualizations directly than we can things in reality). But it goes still further, for at every step the same question will apply. If we cannot know things directly because we can only know our own conceptualizations of them, and we cannot know our own conceptualizations directly because they too must be known through our cognitive faculties, and hence through another conceptualization, then we have set off on an infinite regress of representations of conceptualizations of conceptualizations of conceptualizations, without end. Alternatively, the conceptualist of this stripe could argue that we cannot know things themselves directly because we must know them via our concepts and cognitive faculties, but that we can know our cognitive faculties and concepts directly, and that therefore there is no regress. But to do this requires giving some principled reason why the argument regarding unknowability applies to external things but not to our own concepts and cognitive faculties. It is noteworthy in this connection that Kant was in fact willing to apply his thesis to both external things and our own cognitive faculties. At any rate, it seems clear that it will take an account of knowledge, concepts and reality at least as sophisticated as Kant's to get this version of the "concept" view of information ontologies right, and this before the issue of getting the representation details correct has even been addressed.

### Summary So Far

---

<sup>18</sup> A classic version of this sort of response to the Kantian argument against realism is to be found in the introduction to Lowe, *The Possibility of Metaphysics: Substance, Identity and Time*.

So far we have said that an ontology is a (non-intrinsically-intentional) representational artifact formalized and implemented in a computational system, the intended interpretation of which is what is general in reality, including relations amongst what is general. We have then explored two major answers to the question of what it is that an ontology is primarily intended to represent. The first answer is that an ontology is intended to directly represent general things in *reality*. This answer then subdivides according to the traditional philosophical debates about the question of what the general things in reality are, specifically into *realism*, *nominalism* and perhaps also *conceptualism*. The second answer is that an ontology is intended to represent general things in reality *in virtue of* its representational units (i) being, or (ii) representing *concepts* that in some way (iii) do or (iv) do not represent things in reality. What all of these views have in common is their commitment to the claim that ontologies represent “what is general” (on some interpretation of this phrase) in reality (on some interpretation of the word ‘reality’). The foregoing discussion has thus i) provided a very general criteria that the semantic interpretation of a computational system must meet in order to be considered an ontology, and ii) explored a variety of specific options for how this criteria can be and/or has been met. An example of a definition of “ontology” along these lines is the definition from Smith et. al. that we provided in the previous chapter: “an ontology is a representational artifact whose representational units...are intended to designate universals in reality and the relations between them.”<sup>19</sup> While the authors of this definition acknowledge other possible definitions of ‘ontology’ along conceptualist lines, this definition itself characterizes an ontology as a representational artifact whose representational units represent what is universal directly (not via concepts) and in the realist sense of universals.

#### The representational precision and “logical features” of an ontology

There is a second dimension along which the semantic account of a computational system can qualify it as an ontology, and this is what we have attempted to capture in the second clause of our general definition of an information ontology by saying ‘an information ontology is a representational artifact...which includes a maximal degree of precision regarding the features of what is general and of the relationships between these features.’ By “precision” here we mean to refer to the detail about what is general and the relationships amongst what is general that a

---

<sup>19</sup> Barry Smith, Waclaw Kusnierczyk, Daniel Schober, Werner Ceusters. “Towards a Reference Terminology for Ontology Research and Development in the Biomedical Domain.” Forthcoming in *Proceedings of KR-MED 2006*. ([http://ontology.buffalo.edu/bfo/Terminology\\_for\\_Ontologies.pdf](http://ontology.buffalo.edu/bfo/Terminology_for_Ontologies.pdf))

representational artifact provides in terms of terminological consistency, clear and accurate definitions and statements of identity conditions, rigorously structured taxonomies and partonomies etc.

To see what is at issue here, consider two representational artifacts. The first consists of terms for all of the elements listed in the periodic table of the elements. It includes the information that its terms are to be interpreted as elements of the periodic table (in the realist, nominalist or conceptualist sense), but no other information, either about the meanings of the terms themselves, or about their relationships to one another is included. The second consists of terms for all of the elements in the Periodic Table along with a statement of the intended interpretation of these terms, but includes also scientifically accurate definitions for each of these terms, as well as clear and unambiguous specifications of the relationships in which they stand to each other in the Periodic Table. Now, the first representational artifact amounts to a mere catalogue of *terms*, whereas the second approaches the logical rigor and structuredness of an axiomatic or axiomatizable theory. In order for a representational artifact to be an ontology and possess a ‘maximal degree of precision’ it needs to be much more like the second representational artifact than like the first.

These issues are hinted at by Smith et. al. when they add to the definition cited above the clause “(including a taxonomy of such units)”,<sup>20</sup> as well as later in the same paper; they are also the explicit focus of Guarino and Welty’s definition of an ontology as “a formal mathematical model of some fragment of a conceptualization of a domain.”<sup>21</sup> This definition of an ontology gives pride of place to its being a formal mathematical theory and, at least taken by itself, remains relatively neutral about what it is that the representational units in an ontology should be construed as referring to.<sup>22</sup> Whereas the definition of Smith et. al. focuses on the reality of *what* an ontology represents, the definition of Guarino and Welty focuses on the ontology itself, on formal and logical properties of the *way* that information about what is represented is represented. This is the second dimension along which something can be understood as an ontology in the information science sense.

---

<sup>20</sup> Smith et al., "Towards a Reference Terminology for Ontology Research and Development in the Biomedical Domain."

<sup>21</sup> CITATION

<sup>22</sup> For work towards a definition of ‘ontology’ that takes the formal mathematical features of ontologies very seriously, see Christopher Menzel, "Ontology Theory" (paper presented at the Ontologies and Semantic Interoperability, CEUR Proceedings, vol. 64, 2003).

## Will the Real Ontologies Please Stand Up? Classifying computational systems along the two dimensions:

An information ontology is a representational artifact the intended interpretation of which is what is general in reality, and the relationships amongst what is general, and which includes a maximal degree of precision regarding the features of what is general and of the relationships between these features. Given this definition of information ontology, it should be possible to classify ontologies, on the one hand in terms of what they represent, and on the other hand in terms of the degree of detail or logical structure with which they represent what they represent. Just for example, databases are generally catalogues of particulars, which by this definition excludes them from being ontologies insofar as they are not intended to represent general features of reality. Further, databases do not (at least usually) include a great deal of formal complexity or of explicit computationally tractable information, thus they will not possess a sufficient degree of precision in order to be considered ontologies either. The notions under consideration here, in particular the idea of ‘detail and logical structure’ in an ontology certainly need to be more carefully worked out and defined. Nevertheless, we feel that this basic two-fold characterization of what an ontology is holds out a great deal of promise for distinguishing between ontologies and other kinds of representational artifacts on the one hand, and between different kinds of ontologies on the other.<sup>23</sup>

### Summary So Far:

At this point we have provided a general characterization of what an ontology is and discussed some of the alternative possible ways in which such a characterization has been and can be spelled out. We have also proposed a way in which this characterization might be used for classifying different kinds of representational artifacts and ontologies. All of our discussion so far has focused on crucial features of the semantic interpretation of a computational system that is intended to be an ontology. In the next three sections we will briefly address some important issues at the levels of programming and implementation. A key advantage of thinking about ontologies as computational systems in terms of the tri-level hypothesis is that it holds out the promise of demarcating and helping to structure the work of and conversations amongst the

---

<sup>23</sup> For a presentation that begins to realize some of these ideas, see Stefano Borgo, Nicola Guarino, and Laure Vieu, "Formal Ontology for Semanticists," in *ESSLLI* (Edinburgh: 2005). Slides available on-line at <http://www.loa-cnr.it/Tutorials/ESSLLI1.pdf>.

various interdisciplinary contributors to the project of information ontology, including domain experts, philosophers, logicians, and computer scientists.

### The Nature and limits of the algorithmic level

As stated above, the algorithmic level is the level of programming or software construction and analysis. The question at the algorithmic level is that of what kinds of computer tractable languages and rules to use in order to properly capture the intended interpretation of the information system, and in order to ensure effective and accurate inferences regarding the entities the system is about. It is important to recognize that many different programs or algorithms can be understood as representing the same information and performing the same semantic tasks, and that the same algorithm can be interpreted as representing different information and performing different semantic tasks. There is thus no simple isomorphism between the algorithmic and the semantic levels of an information system, though the two do place certain limitations and constraints on one another. There are, at least, two major issues that arise when considering the algorithmic level. The first has to do with the relationship between the algorithmic level and the semantic level, and tensions that can arise between the desire to represent what is general in reality accurately on the one hand, and to have the information so represented be maximally processable by a computer on the other. The second issue has to do with rendering different ontologies interoperable at the level of algorithms and programming.

Regarding the first issue, results in logic, meta-logic and mathematics from the last hundred years establish that there is a tradeoff between the expressive power and computer tractability of different kinds of logical languages.

#### *Major or Absolute Constraints on the relationship between computation and representation*

- Gödel incompleteness results for axiomatic systems sufficiently strong to express arithmetic and make second-order assertions about reality of the sorts often used and made in the sciences → Capturing some kinds of scientific information in a computational system sufficiently precise to be called an ontology (axiomatized system or near approximation) will involve a loss in automated reasoning possibilities.
- Decidability and undecidability: First-order predicate logic, as well as other highly expressive formal languages, are undecidable. It is not always possible to describe an effectively computable procedure for determining *whether or not* a given conclusion represented in such a language follows from a given set of premises represented in that

same language. If an ontology is explicitly written using an undecidable language, then there will be computational compromises and tradeoffs.

Features and limits of computational systems are not necessarily features and limits of the things that the units in such systems represent (things in reality itself). When we are dealing with natural and biomedical science, we are confronted with the dual necessity of (i) accurately representing domain information in a way that does not compromise the richness, content and import of this information, and (ii) representing domain information in such a way that it can be accessed and utilized by computers for purposes of assisting research, annotation and treatment in the biomedical sciences. The way to handle this, it has been suggested, is to build separate and closely related theoretical and applied information systems called “reference ontologies” and “application ontologies” respectively.<sup>24</sup> The basic idea is to treat a reference ontology as analogous to a scientific theory, focusing in its construction on issues of expressive and theoretical accuracy even at the expense of computational tractability. Application ontologies, then, are software artifacts specifically designed for utilizing sub-portions of reference ontologies to achieve specific computational goals. Application ontologies will, therefore, often make the opposite tradeoff, sacrificing representational detail and accuracy for the sake of computational tractability.

#### *Algorithmic level specific issues*

The second set of major issues to be addressed at the algorithmic level are those that are specific to the domain of computing. These include ensuring the interoperability of different types of programming languages and software, as well as ensuring that such software is maximally effective for purposes of artificial or automated “reasoning”. Even if the information at the semantic level is both accurate and includes a maximal amount of detail and logical structure, this will not be sufficient to make the ontology computer-navigatable or interoperable with other ontologies unless the programs and software in which the ontology is ultimately encoded allow for this.

#### Implementational level

The implementational level, finally, is the level of physical instantiation or hardware. Implementation has to do with the physical device and processes that are used to “run” or

---

<sup>24</sup> Christopher Menzel, "Reference Ontologies / Application Ontologies: Either / or or Both / And?" (paper presented at the Proceedings of the Workshop on Reference Ontologies and Application Ontologies, Hamburg, Germany, 2003).

instantiate the algorithms of a computational system. While most computers run off of silicon chips, many cognitive scientists believe that there is a strong resemblance between the working of computers and the functioning of the human mind, and that therefore neurons are another candidate for the physical implementation of certain kinds of computational systems. For purposes of ontology, at any rate, what is most crucial in understanding a computational system is to understand its semantic and algorithmic components. However, the implementational level also places some basic constraints on what information can be represented, and in what way, insofar as it is in part properties of the physical instantiation of a computational system that will determine its “processing speed” and “memory capacity”. Thus those engaged in ontology design will need to pay some attention to this level of the tri-level hypothesis as well.

### Conclusion

In the foregoing I have suggested thinking of information ontologies as computational systems, and of analyzing them in terms of the Marr-Dawson tri-level hypothesis in terms of the semantic, algorithmic and implementational levels. After suggesting that current information ontologies should not be viewed as instances of artificial intelligence, I proposed the definition of an ontology as a representational artifact (computational system) the intended interpretation of which is what is general in reality, and the relationships amongst what is general, and which includes a maximal degree of precision regarding the features of what is general and of the relationships between these features. This definition gave rise to discussions of universals and of the nature of representation at the semantic level, and to much briefer discussions of formal features of computer programs at the algorithmic level, and to still briefer consideration of the comparatively uncontroversial features and limitations of the implementational level of information ontologies as computational systems. Though I have pushed for a realist understanding of ontology at certain points, my primary goal is that the framework here laid out be able to serve as a basis for discussion amongst advocates of different philosophical interpretations of information ontology, as well as for interdisciplinary conversations amongst philosophers, computer scientists, and practicing natural scientists interested in the construction of information ontologies.

## Bibliography

- Borgo, Stefano, Nicola Guarino, and Laure Vieu. "Formal Ontology for Semanticists." In *ESSLLI*. Edinburgh, 2005.
- Ceusters, Werner, and Barry Smith. "A Realism-Based Approach to the Evolution of Biomedical Ontologies." Place Published, 2006.
- Dawson, M. R. W. *Understanding Cognitive Science*. Oxford: Blackwell Publishers, 1988.
- Johansson, Ingvar. "Bioinformatics and Biological Reality." *Journal of Biomedical Informatics* 39 (2006): 274-87.
- Kelley, David. *The Evidence of the Senses: A Realist Theory of Perception*. Baton Rouge: Louisiana State University Press, 1986.
- Klein, Gunnar O., and Barry Smith. "Concepts Systems and Ontologies: Recommendations Based on Discussions between Realist Philosophers and Iso/Cen Experts Concerning the Standards Addressing "Concepts" And Related Terms." (2005), <http://ontology.buffalo.edu/concepts/ConceptsandOntologies.pdf>.
- Kohler, Jacob, Katherine Munn, Alexander Ruegg, Andre Skusa, and Barry Smith. "Quality Control for Terms and Definitions in Ontologies and Taxonomies." *BMC Bioinformatics*, no. 7 (2006), <http://www.biomedcentral.com/1471-2105/7/212>.
- Kuhn, T. S. "Objectivity, Value-Judgment and Theory Choice." In *Scientific Knowledge: Basic Issues in the Philosophy of Science*, edited by J. A. Kourany. Belmont: Wadsworth Publishing Co, 1986.
- Lowe, E. J. *The Four-Category Ontology : A Metaphysical Foundation for Natural Science*. Oxford: Oxford University Press, 2006.
- . *The Possibility of Metaphysics: Substance, Identity and Time*. Oxford: Oxford University Press, 1998.
- Marr, David. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. New York: W. H. Freeman and Company, 1982.
- Masolo, Claudio, Stefano Borgo, Aldo Gangemi, Nicola Guarino, and Alessandro Oltramari. "Wonderweb Deliverable D18." <http://wonderweb.semanticweb.org/deliverables/documents/D18.pdf>.
- Menzel, Christopher. "Ontology Theory." Paper presented at the Ontologies and Semantic Interoperability, CEUR Proceedings, vol. 64 2003.
- . "Reference Ontologies / Application Ontologies: Either / or or Both / And?" Paper presented at the Proceedings of the Workshop on Reference Ontologies and Application Ontologies, Hamburg, Germany 2003.
- Rayside, Derek, and Gerard T. Campbell. "An Aristotelian Understanding of Object-Oriented Programming." Paper presented at the 15th ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA 2000), Minneapolis, Minnesota USA 2000.
- . "Aristotle and Object-Oriented Programming: Why Modern Students Need Traditional Logic." Paper presented at the 31st ACM/SIGCSE Technical Symposium on Computer Science Education, Austin, Texas 2000.
- Searle, John R. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge: Cambridge University Press, 1983.
- Sider, Theodore. *Four-Dimensionalism: An Ontology of Persistence and Time*. Oxford: Oxford University Press, 2001.

- Smith, Barry. "Beyond Concepts: Ontology as Reality Representation." In *Proceedings of Fois 2004*, edited by Achille Varzi and Laune Vieu, 2004.
- . "From Concepts to Clinical Reality: An Essay on the Benchmarking of Biomedical Terminologies." *Journal of Biomedical Informatics*.
- . "Ontology." In *Blackwell Guide to the Philosophy of Computing and Information*, edited by L. Floridi, 155--66. Oxford: Blackwell, 2003.
- Smith, Barry, and Berit Brogaard. "Sixteen Days." *Journal of Medicine and Philosophy* 28, no. 1 (2003): 45-78.
- Smith, Barry, and Werner Ceusters. "Wüsteria." Paper presented at the The XIXth International Congress of the European Federation of Medical Informatics 2005.
- Smith, Barry, Werner Ceusters, Bert Klagges, Jacob Kohler, Anand Kumar, Jane Lomax, Chris Mungall, Fabian Neuhaus, Alan L Rector, and Cornelius Rosse. "Relations in Biomedical Ontologies." *Genome Biology* no. 6:R46 (2005).
- Smith, Barry, Jacob Kohler, and Anand Kumar. "On the Application of Formal Principles to Life Science Data: A Case Study in the Gene Ontology." In *Database Integration in the Life Sciences (Dils 2004)*, edited by E. Rahm. Berlin: Springer, 2004.
- Smith, Barry, Waclaw Kusnierczyk, Daniel Schober, and Werner Ceusters. "Towards a Reference Terminology for Ontology Research and Development in the Biomedical Domain." Paper presented at the KR-Med 2006.