Ontology for the Twenty First Century: An Introduction with Recommendations

by
Andrew D. Spear

Institute for Formal Ontology and Medical Information Science (IFOMIS), Saarbrücken, Germany
University at Buffalo, Buffalo, New York, U.S.A.

Forward

The material below represents an introduction to ontology, information ontology and ontology design. This introduction is based on the realist approach to ontology and on the associated Basic Formal Ontology (BFO) and principles for best ontology practices that have been developed and researched at IFOMIS in the last four years.

While focusing on work done at IFOMIS, the majority of the discussion is intended to be of a very general nature and to be applicable to a large number of theoretical and applied issues in ontology generally.

The material available here is part of an ongoing project to provide a one-stop resource and introduction for ontology, with special emphasis on principles for best ontology practices. In this spirit we welcome critical comments, suggestions of further resources for inclusion and feedback regarding the accuracy and usefulness of this material.

Acknowledgements

Work on this document has been supported by the Wolfgang Paul Program of the Alexander von Humboldt Foundation, the Volkswagen Foundation Project “Forms of Life” and the Network of Excellence in Semantic Interoperability and Data Mining in Biomedicine of the European Union.

Thanks are due to a number of people for their help with this project:

For detailed comments and feedback, we are grateful to Holger Stenzhorn, Ingvar Johansson, Bastian Fischer and Pierre Grenon.

For valuable conversation and comments, Mathias Brochhausen, Cristian Cocos, Boris Hennig, Kristl Laux, Chris Partridge, Stefan Schulz, Ulf Schwarz, Frederic Tremblay, Barry Smith, and the members of the Institute for Philosophical Studies who were present for the April 11, 2006 meeting in Green Bay, Wisconsin should all be thanked.

Contents

Introduction

A sea of information

In the last fifty years the development of new technologies, the increasing interconnectedness amongst people in all parts of the globe and the continuous increase in the size of scientific, economic and consumer groups, has lead to a veritable explosion in the amount of data that is produced, used and in need of management world-wide. This is especially true in areas such as the biological sciences, medical research and medical practice. In these disciplines thousands of scientists, doctors and clinicians are contributing daily to a massive body of biomedical knowledge and information. Now, like never before in history, the amount of information that is available and being added to daily by the results of new scientific experiments, research and clinical trials constitutes a veritable ocean of extraordinary depth and breadth.

Managing the sea of information

The current and ever-growing stock of biomedical knowledge and information is greater than the knowledge that could realistically and reliably be acquired by any single human individual, even in an entire lifetime. Further, even if it were possible for human individuals to learn or memorize all of currently existing biomedical information, limitations on the speed and power of ordinary human reasoning would prevent such individuals from being able to efficiently retrieve all of the information that they have learned on any given occasion, or to reason about it in systematic ways. Thus, in spite of the fact that massive amounts of data and information are being generated on a daily basis, ensuring the availability of this information to working biomedical researchers and medical practitioners remains a problem. The sea of information is so large and so deep that current researchers and practitioners simply cannot fathom it all. Yet progress in both biological science and medical treatment depends upon current research and treatments making use of the results of previous research and clinical trials. It is thus crucial that the information contained in biomedical textbooks, journals and clinical trial reports be efficiently accessible to and usable by individuals, doctors and research groups other than its original authors, for purposes both of performing research and of treating the actual diseases of actual patients. The problem then, in a nutshell, is to chart the ever-growing sea of information in such a way that its various parts, portions and depths can be efficiently accessed, used, navigated and reasoned about by human individuals.

The obvious solution

The obvious solution to the problem of navigating the sea of information, that is, of managing and making widely available the ever-growing body of biomedical information, has been to make use of the super-human memory and reasoning capacities of computers. In this respect computers have three great virtues. First, they are able to reliably store a tremendous amount of information. Second, they are able to efficiently, reliably and automatically retrieve and reason about the information that they store. Third, information stored in a computerized format can be made instantly accessible to individuals in all parts of the globe via the internet. Thus, if the biomedical knowledge possessed by experts in the various sub-fields of biology and medicine could be organized and stored in interconnected computer repositories, it would be accessible in real time to anyone anywhere in the world, and could be continuously updated in light of new scientific and medical discoveries. Further, the information contained in these databases could, in principle, be used as the basis for certain kinds of automated reasoning that would independently assist in furthering the goals of scientific research and clinical practice. Such a resource would be invaluable for the purposes of both biomedical research and the treatment of patients (imagine a doctor with immediate access to the most current information about all known diseases at the click of a mouse).[1] However, there are a number of obstacles to the achievement of this vision.

Obstacles to the obvious solution

The first obstacle is that the members of the global community of scientific researchers speak different languages, use different terminologies, and format the results of their research in different ways.

The second obstacle is that the computer technology used by these same groups to encode and store their results has, thus far, suffered from many of the same problems. These two obstacles together might be labeled "*the problem of human idiosyncrasy*".

[1] For a discussion of a collaborative effort towards achieving these tasks in the biomedical domain, see Rubin D. L., Lewis Suzanna, Mungall Chris, Misra S., Westerfield M., Ashburner Michael, Sim I., Chute Christopher, Solbrig H., Storey M. A., Smith Barry, Richter J. D., Noy N. F., Musen Mark. "The National Center for Biomedical Ontology: Advancing Biomedicine through Structured Organization of Scientific Knowledge", in: *Omics: A Journal of Integrative Biology*, 10 (2) (2006): 185-198.

A third obstacle is that scientific theories themselves are not merely long lists of true statements about reality, but rather logically interconnected sets of propositions, usually resting on a number of basic concepts and principles. Scientific theories thus have an internal structure and organization that is, as a rule, more detailed and interconnected than are the entries in a standard computer database.

Fourth, the very virtues of computer based implementations (unlimited memory, efficient retrieval and reasoning, widespread availability), in conjunction with the problems already discussed, create a fourth great obstacle to the realization of this ideal. This problem can be summed up by saying that computers are *dumb beasts.* Computers do not understand themselves, their programming or the intended interpretation of the representations that they manipulate. This last fact has a number of implications.[2]

Computers and their discontents

First, unlike human beings, two computers that encode the same information but use different terminology or different organizational principles cannot, for example by sitting down and talking about the matter, come to some kind of agreement or understanding about the common referents of their discrepant terminologies. Rather, the information in two such machines will be separated by a chasm that can be crossed, if at all, only with the help of human intervention. Thus, for two computer-based biomedical information repositories to be compatible or "interoperable", it is essential that they be given either the same terminology, definitions, etc. or an explicit set of instructions for the translation of one terminology into the language of the other and conversely. This might be labeled "*the problem of Babel*".

Second, computers are able to reason only with the information that they have been given. This means that if the original data entered into the computer is vague, ambiguous, contradictory or in some other way unclear, the computer programmed to reason with such data will produce either nothing at all, or else results that are just as vague, ambiguous, contradictory and unclear as the original information it was given, if not more so. This problem could aptly be called "*the problem of nonsense-in-nonsense-out*".

[2] In making this point we do not mean to take a definitive stand on the question of whether or not it is *in principle* possible for a computer or something like a computer to think and reason in the way that human individuals do. We mean only to stress the point that at this time computers, and especially the kinds of databases and information systems being used in the medical and biomedical domains, in fact do not think and understand in the way that human beings do, and that it is this fact about them that leads to many difficulties and problems in realizing the vision of information sharing and use that is under discussion here.

Thirdly and finally, just as a computer can only reason with what it has been given, so it can only represent what it has been given as this has been given. Computers cannot check the information that they contain for factual accuracy to reality. This means that the information contained in a computer database will only be as accurate to the facts of a given scientific or practical domain as the persons constructing this database have taken the time to ensure it to be. This might be called "*the problem of computer-information-solipsism*".

The shape of solutions sketched

There are thus a number of problems in need of a solution before the use of computers to store, manage and reason about complex biomedical information can be expected to have wide-spread and optimal success. The idiosyncrasies of both biomedical researchers and practitioners, as well as of computer programmers and data-base builders need to be overcome. Computer languages and programs with expressive power sufficient to handle the richness of scientific theories need to be developed and implemented. Substantial care needs to be taken at the outset to ensure that the terminology, definitions, etc. that are entered into biomedical information databases are interoperable between databases, internally coherent and well-defined, and accurate to the facts of reality as reflected in the current (and developing) state of knowledge possessed in the biomedical sciences. The basic answer that has been proposed and is being developed as a general response to all of these problems is "ontology".[3]

Ontology

Traditionally, ontology has been defined as the philosophical study of what exists: the study of the kinds of entities in reality, and the relationships that these entities bear to one another.[4] While this study includes the entities dealt with by the specialized sciences (physics, chemistry, biology, etc.), it also has a more general focus, one directed at providing an account of the most general or basic features of reality: an account of the kinds of objects and relations that are common to all scientific domains whatsoever. Examples of such general or common features of reality might include identity, both at a time and across time, qualities such as color

[3] For a general and recent discussion of many of these points, and for what might be called a summary of many of the things to come in this manual, see Barry Smith and Werner Ceusters. "Ontology as the Core Discipline of Biomedical Informatics: Legacies of the Past and Recommendations for the Future Direction of Research", in: Dodig Crnkovic Gordana, Stuart Susan (eds.). *Computing, Philosophy, and Cognitive Science.* Cambridge Scholars Press, Cambridge, forthcoming (http://ontology.buffalo.edu/medo/Recommendations_2005.pdf).

[4] For more on this see Barry Smith. "Ontology." in L. Floridi (ed.), *Blackwell Guide to the Philosophy of Computing and Information*, Oxford: Blackwell, 2003, 155–166.

and shape, compositional structure involving relationships of part to whole, relationships such as causation and class-membership, etc. Thus, whereas the biologist studies, for example, cells and the physicist studies, for example, atoms, the philosophical ontologist is interested not only in studying (usually with the help of the natural scientist) these entities themselves, but primarily in giving an account of what is common to both cells and atoms, and of the relationships in which these kinds of entities stand to one another, relationships that may well extend across the normal disciplinary boundaries of specialized sciences such as biology and physics. The goal of philosophical ontology is to provide clear, coherent and rigorously worked out accounts of the basic structures to be found in reality.[5]

In recent times use of the term ‘ontology’ has become prominent in the areas of computer science and information science research and in the applications of these fields to the management of scientific and other kinds of information. Here ‘ontology’ has the meaning of a standardized or agreed upon terminological framework, of varying generality, in terms of which information repositories of different sorts are to be constructed. The purpose of such ontologies is both to give an articulate internal structure to electronic information repositories, and to make possible the interoperability or inter-translatability of different repositories containing different information, in such a way that the information in both repositories can be understood in terms of a common language. The ontological problem of information repository construction and management is not, however, simply the problem of agreeing on the use of a common vocabulary. Rather, it is the problem of adopting a (sometimes very general) set of basic categories of objects, of determining what kinds of entities fall within each of these categories of objects, and of determining what relationships hold within and amongst the different categories in the ontology. The ontological problem for computer and information science is thus identical to many of the problems of philosophical ontology, and it is becoming more and more clear that success in the former will be achievable, if at all, only by appeal to the methods, insights and theories of the latter. [6]

How ontology can help

[5] For more detailed discussion of ontology and of related issues in metaphysics, see E. J. Lowe, *A Survey of Metaphysics,* Oxford: Oxford University Press, 2002 & *The Four Category Ontology,* Oxford: Oxford University Press, 2006.

[6] For more on ‘ontology’ in the philosophical and in the information scientific use, see Barry Smith. “Ontology”, and also Douglas Mayhew and Dirk Siebert. "Ontology: The Discipline and the Tool.” in: Büchel Gregor, Klein Bertin, Roth-Berghofer Thomas (eds.). *Proceedings of the First Workshop on Philosophy and Informatics, Deutsches Forschungszentrum für künstliche Intelligenz*, Cologne, 2004, 57-64.

The ontological approach to solving the problem of computer based biomedical information management addresses the four major obstacles discussed above in roughly reverse order. The ontological approach begins with human researchers attempting to specify in detail the information that is to be computer-implemented. This includes selecting appropriate terminology, defining this terminology in a rigorous, clear and logically coherent fashion, and ensuring, as much as possible, that the information to be implemented is accurate to the facts of reality. Careful attention to these issues ensures that the internal structure of the scientific theory of interest is maintained and implemented. Further, structuring the information to be implemented in terms of general ontological theories, theories about relationships and objects that are common to different scientific domains, helps address the problems of interoperability or "communication" between different computer artifacts.[7]

Also, beginning with a rigorous and logically coherent specification of the theoretical information to be implemented makes it possible to address the problems of human idiosyncrasy. If different groups of researchers use different terminology for the same entities and to express the same scientific truths, it should be possible to inquire into the reasons for this. If the difference in terminology and expression is rooted in theoretical differences, these differences can be made explicit and discussed, in a scientific fashion. If the differences are merely terminological, then in the interest of making scientific discourse more intelligible, it should be possible to get the two groups to accept a single standardized term or set of terminology, at least for the purposes of interacting with common repositories of information for their scientific field.

Finally, the issue of computational idiosyncrasy can be addressed. Given a formal specification of the theoretical information to be encoded, along with a human consensus about its terminology and a solid understanding of the internal structure of the theory, the only questions of computation that are relevant are first that the program or programming language that is selected have a sufficient amount of expressive power to include all pertinent information of the theory, second that it have sufficient reasoning power to support automated information retrieval and integration, and third that it be interoperable with other such systems. In short, once it is clear what the content and goals of the meaningful scientific theories involved are, computational idiosyncrasy should largely disappear.

---

[7] Such general ontologies are often called "formal" or "top-level" ontologies, and more will be said about them in the sections to follow.

In what follows: towards an *Organon* for the Information Ontology Age

In the following pages the basic elements of the ontological approach to solving the problems of information management will be presented and explained. In addition, specific recommendations and principles will be put forward. These are intended for use by individuals interested in the possibility of constructing information ontologies of their own.

This discussion is intended to take place specifically against the background of the use of computer based ontologies for the purposes of science, scientific research, and the application of scientific results to particular problems, as in medical practice. In the age of computers there are many potential uses for terminologies and ontologies in the organization of electronic information. However, we are here concerned specifically with the application of ontology in the domain of science. Taking science to be the systematic attempt to account for and explain what *actually* exists, its behavior and its nature, our primary focus will be on the construction of ontologies for the purpose of representing what exists, not for keeping track of fictional discourse, daydreams or mythological creatures.

This document is thus written in the spirit of an *Organon,* an instrument for the proper conduct and representation of scientific research. The first *Organon* was written by the Ancient Greek philosopher Aristotle in the 4th Century B.C., and included his works on logic and the theory of science.[8] The second great *Organon,* the *Novum Organum* (1620) of the Englishman Francis Bacon was written as an update, extension and correction of the Aristotelian *Organon* in light of the success and experimental methods of modern natural science almost 2000 years latter.[9] All sciences have some methodological assumptions in common; the principles of logical inference and reasoning are one example, while the complicated structure of scientific verification and falsification, and adherence to principles of theory evaluation such as "simplicity", "explanatory power" and "predictive power" are others. It is possible for any given scientist or scientific community to employ these methods with greater or lesser adequacy, and the goal of an *Organon* is to codify, so far as is possible, the principles for best or most adequate scientific research and representation practices.

The dawn of the computer age, and the increasing use of computers in scientific research, both for representing information, and for acquiring new results, has raised a host of new

[8] The works known as Aristotle's *Organon* can be found in *The Complete Works of Aristotle*, Two Volumes (Jonathan Barnes ed.). Princeton: Princeton University Press, 1984.

[9] Francis Bacon. *Novum Organum* (Urback, P. and Gibson, J. transl. and eds.). Chicago: Open Court, 1994.

questions about the nature of science, the nature and limits of computability and the nature of representation itself. Just as the advent of Modern Science during the Renaissance made it necessary for Francis Bacon to update the *Organon* of Aristotle with a new organum that would take into account the scientific developments of his time, so what is needed at the beginning of the 21st century is another new organum, an "*Organon* for the Information Ontology Age", one that will clarify the basic principles and methods of computer-based and computer-assisted scientific representation and research. This document is not such an *Organon*, but is intended as contribution towards the realization of that task.

Chapter 1:
What does it mean to "build an ontology"?

The questions "what is an ontology" and "how does one use the ontological approach to organize the sea of information" are closely related, and it will become ever clearer as we move forward that answering the second question involves answering the first, and vice versa. The purpose of this section is to provide an initial overview of the answers to these two questions, one that will begin to establish the framework for understanding and applying the specific methods and principles that will be discussed in sections to come.

The following definition of 'ontology' has recently been proposed, and it contains most of the elements that it will be important to discuss here: an ontology is a representational artifact whose representational units are intended to designate universals in reality and the relations between them.[10]

This definition has two parts. The first identifies an ontology as a representational artifact consisting of representational units, while the second has to do with what the representational units in such an artifact are intended to refer to or be about. We will deal with each of these in turn.

Representations

The human world is full of representations and representational artifacts. The key feature common to all representations is that they make reference to or are about something else. Thus a representation is an idea, an image or a description that refers to some entity or entities external to itself. The memory that I have of the Tower Bridge in London is a representation in my mind

[10] 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).

that is about or refers to an entity other than itself, namely the actual Tower Bridge that exists on the Thames River in London. Similarly, the thoughts of a scientist as she looks through a microscope at bacteria, namely the thoughts that "these are bacteria", are mental representations that, taken together, point beyond themselves and make reference to the actual existing bacteria that are under investigation. It is, indeed, one of the most basic features of human thought that beliefs, desires and experiences in general point beyond themselves and refer to the objects that they are about.[11] However, representations by themselves are not yet ontologies in the sense in which we are here interested. Ontologies have the important further feature of being representational artifacts.

Representational artifacts

A representational artifact is an entity which makes publicly available pre-existing cognitive representations from the minds of its author or authors. Representational artifacts include things such as signs, books, pictures and diagrams. A key feature of representational artifacts it that they include ledgers or rules for their interpretation. Thus, maps do not simply come color coded, they also come with a key or table that makes it possible to interpet their color coding as representing certain kinds of things (countries, oceans, mountain ranges, etc.), and the words in which these tables and keys are written themselves have publicly available rules for their interpretation as referring to things in the world, namely the semantics of natural language itself.

A simple kind of representational artifact would be if I were to draw a picture of the Tower Bridge in London based on the mental representation that comprises my memory of having once seen it.[12] Whereas my memory is a cognitive representation, the picture that I draw based on it is a representational artifact. I intend for it to refer to the same thing to which my original memory refers to, namely the Tower Bridge, and yet it exists independently of my own mind or thoughts in a form that is publicly observable and inspectable.

---

[11] It is important to note that representations need not always be mental or cognitive in nature. A picture of the Tower Bridge in London is also a representation, something that makes reference to an entity other than itself. For more on the directedness or aboutness of mental states in general, and on the distinction between mental and non-mental representations in particular, see John R. Searle. *Intentionality: An essay in the philosophy of mind.* Cambridge: Cambridge University Press, 1993. See especially Searle's distinction between "derived" and non-derived or original intentionality.

[12] It should be noted that everything to be said in what follows would hold good also if, instead of a picture of the Tower Bridge, I produced a detailed handwritten description of it. This document would also be a representational artifact, and all that will be said about the picture of the Tower Bridge would apply equally well to a written description.

Note here that just as my memory of the Tower Bridge can be better or worse, more or less accurate, so also the representational artifact that I create based on this memory can be better or worse, that is, more or less accurate as a representation of the real entity it is intended to refer to. Notice also that once I render my mental representation into the form of a representational artifact, it becomes available to the community at large for inspection and use. If the drawing (representational artifact) I have created is a good one, someone who was previously unfamiliar with the Tower Bridge may gain new knowledge about an object that exists in the world from inspecting it, while if the representation I have created is a bad one, someone else who has also seen the Tower Bridge will be able to criticize it and suggest needed improvements. There are two points of crucial importance to note here.

<u>Representational artifacts normally represent things, not mental representations, concepts or memories</u>

First, when I attempt to create a representational artifact that makes reference to the Tower Bridge by drawing a picture, it is not the mental representation in my head, the actual memory of the Tower Bridge, that I am trying to draw, but rather, the Tower Bridge itself. This is a very important point to keep in mind when constructing an ontology, and there are a number of ways to see that it is true. First, what would it mean for me to attempt to draw a picture (create a representational artifact) *of my memory*, that is, of my mental representation of the Tower Bridge? What would such a picture look like? It is quite arguable that no one really knows what a drawing of a memory would look like. But more importantly for the purposes of ontology, the answer to this question is unimportant.[13]

In creating a drawing of the Tower Bridge (representational artifact) based on my memory of having seen it (cognitive representation), my primary concern is with accurately representing, not my memory of the Bridge, but the Bridge itself (the thing in reality). Should I have an opportunity to see the bridge again in the future, and to compare it with the drawing that I have made, I may well identify a mistake or an absence of detail in the drawing and decide to correct it in order to more accurately represent the bridge, even if my original memory of the bridge contains no such additional information. Additionally, if other people look at my drawing

[13] Except, perhaps to an empirical psychologist. But note that, even in the case of empirical psychology, where researchers might actually be interested in representing facts about memories or concepts as scientific objects, the goal of these researchers will not be to represent their memory of memories or their concepts of concepts, but rather to represent the actually existing neuro-physiological/mental objects in reality that are the objects of empirical psychological research.

of the Tower Bridge and criticize it, they will engage in this criticism by citing facts, not about my memory or my cognitive representations, but about the drawing and about the bridge itself. And, as far as ontology for the domain of natural science is concerned, this holds true in every case. When constructing a representational artifact for use in science, such as an ontology, based on cognitive representations or concepts in the minds of individual subjects, the goal is not to accurately represent in a publicly accessible way the representations or concepts that exist in the individual's minds, but rather the things in reality that these representations are representations of.[14]

It is possible both to *use* and to *mention* representational artifacts, but this distinction must be respected

The second important point to be made in connection with representational artifacts is that there is a fundamental distinction between *using* such artifacts to make reference to things in reality, the entities that they are representations of, on the one hand, and *mentioning* such artifacts by engaging in discourse about them on the other. The construction of coherent functional ontologies requires that this *use-mention* distinction be strictly taken into account. Consider first an example from ordinary language. One can use the words 'The Tower Bridge' to refer to an object in reality, as in "The Tower Bridge is a well-known structure on the Thames River in London". However, one can also mention the words 'The Tower Bridge', as in " 'The Tower Bridge' is a set of words used primarily by speakers of English to refer to a structure on the Thames River in London, and these words are made up of fourteen occurrences of letters from the Latin Alphabet".

Similarly, to return to the example of a drawing of the Tower Bridge, one could use such a drawing in order to explain to someone what the Tower Bridge is, and what its characteristic features are. In this case the representational artifact is being used in order to talk about the

[14] 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", in *Journal of Biomedical Informatics*, forthcoming, Barry Smith. "Beyond Concepts: Ontology as Reality Representation" In Achille Varzi and Laure Vieu (eds.), *Proceedings of FOIS 2004. International Conference on Formal Ontology and Information Systems,* Turin, 4—6 November 2004, Gunnar O. Klein & 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." (http://ontology.buffalo.edu/concepts/ConceptsandOntologies.pdf), Ingvar Johansson. "Bioinformatics and biological reality." *Journal of Biomedical Informatics* 39 (2006) 274—287, Barry Smith and Werner Ceusters. "Wüsteria", in *Proceedings of the XIXth International Congress of the European Federation for Medical Informatics* (MIE 2005), (http://ontology.buffalo.edu/concepts/ConceptsandOntologies.pdf).

object that it makes reference to. But one could also mention such a drawing by making it and its properties the explicit theme of discourse and saying things like "this drawing is made of paper", "this drawing was done with a mixture of pencil and pastel" and "this drawing leaves important details of the Tower Bridge out of account". In all such cases discourse *mentions* the representation and is thus explicitly about the representation itself, not about the objects or entities to which the representation refers. It is a very common mistake in the construction of ontologies to mix statements that are meant to be part of the ontology, and thus refer to objects in reality on the one hand, with statements that are about the ontology (representational artifact) itself, and thus are meant to refer to items, entries and components within the ontology, but not to the representation-independent things in reality. Some examples include the following:

In the Medical Subject Headings (Mesh) can be found the following hierarchical relationship "National Socialism *is_a* MeSH Descriptor".[15] Here National Socialism, which is a kind of political movement that existed in the world, is identified as a kind of term in the MeSH database, however the definition of 'National Socialism' in MeSH as "The doctrines and policies of the Nazis or the National Social German Workers party, which ruled Germany under Adolf Hitler from 1933-1945. These doctrines and policies included racist nationalism, expansionism, and state control of the economy. (from Columbia Encyclopedia, 6th ed. and American Heritage College Dictionary, 3d ed.)" makes it clear that the use and mention of the term 'National Socialism' have here been confused. Similarly, in the National Cancer Institute Thesaurus (NCIT), the following definition of 'Conceptual Entities' can be found "An organizational header for concepts representing mostly abstract entities".[16] As a definition there are a number of problems with this. However, what is important here is simply that, once again, the use and mention of a term have been conflated. Whereas we would expect a database that is about things in reality to provide definitions that would tell us facts about the basic features of those things (in this case the basic features of conceptual entities), what we get here is an explanation of how the term 'conceptual entity' is used as a part of the representational artifact that is the NCIT. Once again, statements about things (use) and statements about the words used to mention things (mention) have been conflated. Other examples include the definition of 'mouse' as "name for

[15] MeSH accessed via UMLS Knowledge Source Server Release 2006AC. http://umlsks.nlm.nih.gov/kss/servlet/Turbine/template/admin%2Cuser%2CKSS_login.vm;jsessionid=421388A81A91D5EC4C8759E7D5B3B13A.kss2, Thursday, September 28, 2006.
[16] NCI Thesaurus accessed Via UMLS Knowledge Source Server, Release 2006AC, Thursday, September 28, 2006.

the species *mus musculus*" in BIRNLex,[17] the entry "Bacterium causes Experimental Model of Disease" in the Unified Medical Language System (UMLS),[18] the definition of 'animal' as "a subtype of Living Subject representing any animal-of-interest to the Personnel Management domain" from the HL7 Glossary, [19] and "living subject is_a code system" from the HL7.[20] All of these are examples of ways that databases and ontologies can go wrong when they fail to keep separate the use of the ontology and its terms to refer to things in reality, and the mention of terms and elements of the ontology in discourse that is explicitly about it, its construction and its constituent elements.

The example of a representational artifact that we have focused on so far, a drawing of the Tower Bridge, is admittedly quite simple. More sophisticated representational artifacts include certain aspects of natural language, maps, blue-prints, and various classificatory schemes employed by the natural sciences, notably the periodic table of the elements. However, it is important to recognize that the basic distinction between representing a memory or concept and representing a thing, and between using a representational artifact and mentioning it, apply in these more sophisticated cases as well. It is one thing to talk about (mention) the Periodic Table of the Elements as an important development in the history of human knowledge, and it is quite another thing to use this same Table in order to learn something about the nature of elements in reality. What all of the things mentioned above have in common is that they are codified publicly available representations possessing interpretive rules for understanding how they refer to reality, which can be publicly evaluated for accuracy and usefulness.

To say, then, that an ontology is a representational artifact is to say that it is an explicit and publicly available rendering of a representation or system of representations belonging to an individual or group of individuals. Paradigm examples of ontologies in this sense would be scientific textbooks (or at the least, portions thereof), geographical maps, and certain kinds of information databases. When such a representational artifact is formalized, that is, when such an

---

[17] UMLS Semantic Network http://semanticnetwork.nlm.nih.gov/

[18] UMLS Semantic Network http://semanticnetwork.nlm.nih.gov/, Also to be found in the UMLS are "Experimental Model of Disease **affects** Fungus", "Experimental model of disease **is_a** Pathologic Function", "cancer documentation is_a cancer", "living subject is_a information object representing an animal or complex organism", "individual allele is_a act of observation".

[19] Various Contributors eds., HL7 Publishing Technical Committee. Last Published 11/22/2005 8:05 PM. HL7® Version 3 Standard, © 2005 Health Level Seven®, Inc.

[20] HL7 Version 3.0 accessed Via UMLS Knowledge Source Server, Release 2006AC, Thursday, September 28, 2006.

artifact is expressed in a logical or programming language of some sort, it can be called a 'formalized representational artifact'. Formalized representational artifacts have the advantage, normally, of being both rigorously formulated and computer implementable. It is thus ontologies in the sense of formalized representational artifacts that currently hold out the greatest promise of enabling computers to help human researchers and clinicians cope with the ever-growing sea of biomedical information.

Ontological reality, an initial introduction

It is the second part of the above definition, the point that the representational units of an ontology "are intended to designate universals in reality and the relations between them," that is crucial to an initial understanding of what an ontology is. From what has been said so far, it might be thought that any representational artifact at all is or could be considered an ontology. This is false. While there are many kinds of representational artifacts, a representational artifact is an ontology only if the intended referents of its representational units are real universals and real relationships amongst such universals on the side of reality.[21] But what is a universal, and what is meant by talk of relationships amongst universals?

It is a basic assumption of scientific inquiry that nature is structured, ordered and regular, at least to some degree. Though scientists always perform experiments and make observations regarding particular objects, what they are actually interested in are the generalizations about the structure, order and regularity that exists in nature that such experiments and observations make possible. Universals are that which is general or abstract in reality. They are the philosophical explanation of the structure, order and regularity that is to be found in nature, and they are what all members of a natural kind, grouping or species (for example the kind "feline" or "mammal") have in common. Universals are repeatable in the sense that they can be instantiated by more than one object and at more than one time, whereas particulars, such as myself, Tibbles the cat and specific political administrations, are non-repeatable, they can exist only in one place and during one period of time. Because of this, universals do not have a determinate location in space or time. Rather, they exist at all times and in all places where particular entities instantiating them exist.

[21] In this connection it should be noted that while the picture of the Tower Bridge discussed above is indeed a representational artifact, it is not an ontology in the strict sense, since it represents only a single particular thing in reality and not, for example, information about bridges in general.

One way of characterizing a universal is to say that it is that in virtue of which a thing is what it is, and without which that thing would not be the kind of thing that it is. For example, Aristotle, a well-known Ancient Greek philosopher, believed that the universal "human" is characterized by the features of being an "animal" and being "rational". Given this characterization of the universal "human", the following statements should also all be true:

- all humans are rational animals;
- it is in virtue of being a rational animal that any given entity is a human;
- to understand what a rational animal is is to understand what it is for a entity, any entity, to be human; and
- any entity that lacks one or both of these qualities is not a human being;
- the universal "human" has existed, exists now, and will exist in the future, regardless of the knowing activities of scientists, it neither came into existence when Aristotle first thought about it, nor will it go out of existence if everyone stops thinking about it.

Whether or not Aristotle was entirely correct about the nature of human beings is not important here. What is important is that to know the nature of a thing, for example a human being or a chemical element, is to know the nature of the universal(s) that that thing instantiates and that, taken together, comprise the nature of that thing. Thus, the goal of scientific research in any given domain is to discover the nature of the universals that are instantiated by entities in that domain, whether in biology, chemistry, physics, or some other science.

As opposed to universals, particulars are the individual denizens of reality. Particulars instantiate universals, but cannot themselves be instantiated. It is in virtue of instantiating a given universal that two particulars will be similar in some respect. For example, to say of two entities that they are both cats, and that they are similar in virtue of being cats, is to say that they both instantiate the universal "cat". Particulars thus exist in space and time, and come into and pass out of existence. It is possible to causally interact with, directly see with one's eyes, touch and smell particulars, but not universals. For example, the universal cucumber is an abstract entity that is instantiated by and accounts for the similarities amongst all particular cucumbers, but unlike its instances, the universal has never been and cannot be sliced, diced, made into salad or consumed with a cool glass of chardonnay on a warm summer afternoon.

Whereas particulars are fleeting and contingent, universals are what is abiding and permanent in reality. To say that an ontology is concerned with representing universals on the side of reality is thus to say that an ontology is a representational artifact the primary purpose of which is to represent what is essential, law-like and general in reality. The way to begin doing this is to represent, as accurately as possible, the general features that are attributed to reality by the natural sciences, such as biology, chemistry and physics. It is important to distinguish, however, between genuine universals on the one hand and mere classes on the other.[22]

Every universal has a corresponding class, but not every class corresponds to a universal. A class can be defined as a collection of particulars falling under a term in such a way that the term applies to every member of the collection, and every particular to which the term applies is a member of the collection. For example, the class corresponding to the universal "cat" will be designated by the term 'cat' and will contain all and only the particular cats that exist in reality. However, there are many classes that do not correspond to any universal, and these fall into two general kinds.

The first are classes designated by arbitrary general terms. One example of an arbitrary general term would be the term 'grue', which was created by the philosopher Nelson Goodman to pick out all and only things that are green before a certain time (say the World Cup Tournament of 2006), and then blue at all later times[23]. The problem with this class, and the reason that it does not correspond to any universal, is that there is not any general feature of reality having to do with the possession by an entity of a certain color, such as green or blue, relative to its existing before or after a certain time alone. Thus, even if there are individuals in reality that fall within the extension of this bizarre term, it is highly unlikely that they do so in virtue of instantiating some important common nature or universal. Thus 'grue' may well define a class, but it does not correspond to a universal on the side of reality.

The second kinds of classes that do not correspond to universals in reality are classes created by using a general term to make reference to particulars existing at a specific time or in a specific place, such as the class of all women currently living on the north coast of Germany, the class of all athletes over the age of 30, or the class of all individuals currently infected by HIV on the Continent of Africa. There may be important reasons to want to talk about classes such as these, but it is equally important for ontological purposes to recognize that these classes do not correspond to universals. The reason for this is that these classes make explicit reference, not to

[22] For more on this, see Barry Smith, Waclaw Kusnierczyk, Daniel Schober, Werner Ceusters. "Towards a Reference Terminology for Ontology Research and Development in the Biomedical Domain."

[23] Nelson Goodman. "The New Riddle of Induction", in *Journal of Philosophy* 63 (1966): 281-331.

general kinds, but to collections or groups of particular entities. Whereas the universal "Oxygen" does not exist at any particular time or place, and is a general law-like feature of reality, the class defined by the expression 'all of the oxygen in this jar' refers only to a particular collection of particular oxygen molecules in a specific location and at a specific time. These particular molecules exist at a certain time and in a certain place, they have a specific temperature and are mixed with a certain ratio of molecules of other sorts, however none of these things is true of the universal "oxygen" (additionally, it is possible for a person to breathe the oxygen molecules in the jar, but it is not possible for any person to breathe the universal "oxygen").

Some examples of arbitrarily defined classes from the International Classification of Disease (ICD 10) include: "other problems with special functions", "tuberculosis of unspecified bones and joints", "tubercle bacilli not found by bacteriological or histological examination, but tuberculosis confirmed by other methods (inoculation of animals)", "other mineral salts, not elsewhere classified, causing adverse effects in therapeutic use".[24]

Universals, classes designated by arbitrarily defined general terms, and classes defined in terms of groups of particulars existing at particular times and places are all important for various purposes. However, for the purposes of constructing an ontology it is essential that the classes designated by each of these things be kept separate, and that primary importance be given to specifying accurate representations of universals. More will be said about these issues in subsequent sections.

The other important notion mentioned in the definition of an ontology that we have given is that of "relationships" holding amongst universals in reality. The general idea of a relationship is familiar from both everyday language and the theories of science. If John and Mary are spouses, then they stand in the relationship of being married (and perhaps also of being in love). If an event such as the collapse of a bridge occurs immediately after some other event, such as an explosion directly underneath the bridge, we say that the explosion stands to the bridge-collapse

[24] International Classification of Disease 2006 Version, http://www3.who.int/icd/currentversion/fr-icd.htm, accessed Thursday, September 28, 2006. Further examples include: "other general medical examination for administrative purposes", "assault by other specified means", "other accidental submersion or drowning in water transport accident injuring other specified person", "accident to powered aircraft, other and unspecified, injuring occupant of military aircraft, any rank", "other accidental submersion or drowning in water transport accident injuring occupant of other watercraft – crew", "normal pregnancy", "fall on stairs or ladders in water transport injuring occupant of small boat, unpowered", "railway accident involving collision with rolling stock and injuring pedal cyclist", "injury due to war operations by lasers", "nontraffic accident involving motor-driven snow vehicle injuring pedestrian".

in the relationship of cause to effect: the explosion causes the bridge to collapse. For our purposes here, what is important to recognize is that it is not only universals, but also the relationships that exist amongst them (themselves a kind of universal) that comprise the contents of scientific knowledge. It is one thing to know something about the species "feline", it is another and much better thing to know also how the species "feline" fits into the larger picture of living things in nature, in particular, what its relationship is to other species. Similarly, it is one thing to understand something about the universal "hydrogen", it is another and much better thing to know how "hydrogen" is related to other element-universals, relations that are captured, for example, in the Periodic Table of the Elements. Often, full knowledge of a given universal requires understanding also the relationships in which it stands to other universals of similar kinds, and conversely. It is for this reason that ontology concerns itself not only with representing universals, but also with representing explicitly the relationships that obtain amongst universals. It is only by doing this that the full internal structure and content of scientific theories can be accurately and completely represented.[25]

As an initial and very general clarification of what is meant by talking of the relations amongst universals that an ontology represents, it is important to distinguish between three general kinds of relationships. Relationships that hold between two universals, relationships that hold between a universal and a particular, and relationships that hold between two particulars.[26]

The paradigm example of a relationship that holds between universals is the "is_a" relationship, as in "feline is_a mammal". The relationship of "is_a" holds amongst universals in virtue of the fact that they stand in hierarchies of generality. For example, the hierarchy extending from the universal "feline" through the universals "mammal", "animal" and "organism" can be understood as structured from least to most general in terms of the "is_a" relation. Thus, more concrete universals stand in "is_a" relations to, or are "subsumed_by" more general universals.

---

[25] It is important to note here that there is a distinction between relations that are particulars or "instance-level" relations, and relations that are universals, just as there is a distinction between particulars and universals more generally. We will not address this issue in any great detail here, but the interested reader is referred to Mertz, D. W. *Moderate Realism and its Logic*. London: Yale University Press, 1996, for a detailed discussion of these matters.

[26] For more on this, see Barry Smith, 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* 6, no. 6:R46 (2005).

A paradigm example of a relationship between a particular and a universal is the "instantiates" relation, as in "Tibbles instantiates Tabby", where Tibbles is a particular flesh and blood cat, and "Tabby" is the universal for the feline breed of Tabby cats. All particulars stand in the relation of "instantiates" to some universals, but no universal can properly be said to instantiate any particular. Further, no particular stands in an "is_a" relationship to any universal.

Finally, a paradigm example of a relationship holding between particulars is the "part_of" relation, as in "John's left leg part_of John". The relationship "part_of" is a pervasive one, and is one that is very important for many kinds of biomedical science. There are many other relationships that hold amongst particulars, such as "is_cause_of" and "depends_on", and it is important to note that, under specific conditions that will be discussed in more detail later, such relationships can also be construed as obtaining amongst universals.

More will be said about relations in the following sections. For now, what is important to recognize is that there are different kinds of relationships, some of which hold amongst universals, and that fully understanding a given scientific domain requires not only knowing what universals exist in that domain, but also what kinds of relationships hold amongst those universals.

The definition of 'ontology' reconsidered

Given this clarification of the nature of universals and the relationships that obtain amongst them, the definition of an ontology as "a representational artifact whose representational units are intended to designate universals in reality and the relations between them," should be essentially clear. The basic idea is that ontologies are about what is general, structured and law-like in reality. Further, ontologies represent this generality not only by containing general or common terms representing universals, but also by capturing and explicitly representing the relationships that obtain amongst these universals. This differentiates ontologies on the one hand from *terminologies,* representational artifacts containing natural language terms and definitions, but not rendering explicit the structure or relationships amongst the entities referred to by these terms (such as the ULMS), and on the other from *inventories*, representational artifacts designed to keep track of particulars, of what is specific in reality.

What does it mean to construct an ontology?

When one sets out to construct an ontology then, what one is doing is designing a representational artifact that is intended to represent the universals and relations amongst universals that exist, either in a given domain of reality (such as that studied by molecular biology), or across such domains.

Chapter 2: Ontology and other things: the nature and uses of information ontology

This section, still in progress as of now, will provide a general theoretical framework and definition incorporating both realist and conceptualist approaches to ontology. It will then make some general statements applicable to all information ontologies whatsoever, and discuss the differences between information ontologies on the one hand, and things like terminologies, thesauri, taxonomies and etc. on the other. Finally, consideration will be given to some of the specific actual and possible uses of information ontologies.

Some information about the topics of this chapter can be found in the following documents and slides-

- Daniel Schober. "Recommendations." http://msi-ontology.sourceforge.net/recommendations/.
- Stefano Borgo, Nicola Guarino, Laure Vieu "Formal Ontology for Semanticists", http://www.loa-cnr.it/Tutorials/ESSLLI1.pdf

Chapter 3: What are the methods of ontology?

In the introduction, we articulated the problem of managing the sea of biomedical and scientific information, and discussed ontology as the proposed general solution to this problem. In the last sections, an ontology was defined as a representational artifact whose representational units are intended to designate universals in reality and the relations between them; and it was noted that it is when such representational artifacts are regimented and formalized in certain computer tractable ways that they begin to genuinely contribute to the resolution of the "sea of information" problem. Constructing an ontology is the problem of constructing a formalized representational artifact whose representational units designate universals in reality and the relations between them; but what does this process look like? This section sketches an initial outline.

The ontological solution to the problem of the sea of information is essentially a top down one. This means that it begins with theoretical considerations of a very general nature, making the assumption that keeping track of very specific information about organs, genes and diseases requires getting the very general scientific and philosophical details behind this information right, and doing so in a systematic and coherent fashion. It is only once this has been done that the detailed knowledge base of the biomedical sciences can be encoded in such a way as to ensure widespread accessibility and usability. Importantly however, regimenting all of this information in such a way that it can be kept track of and reasoned about by computers requires, for reasons discussed in the introduction, that a maximum amount of clarity and precision be used at each step in the process of identifying and defining universals and the relations amongst them in a given domain. The general method to be followed in constructing an ontology can be summarized in the following steps:

- Explicitly determine and demarcate the subject-matter or domain of the ontology.
- Gather information: Determine what the universals and relations amongst universals dealt with in this subject-matter are.
- Concretize this information in the form of a representational artifact, such as a written document, grid, etc.
- Regiment the information contained in this representational artifact in order to ensure:
  i. Logical, philosophical and scientific coherence,
  ii. Coherence and compatibility with other relevant ontologies, and
  iii. Human intelligibility.
- Formalize the regimented representational artifact in a computer tractable language.
- Implement the representational artifact in some specific computing context.[27]

Determining the subject-matter of an ontology has a number of components. Important among these are dealing with the distinction between formal and material ontologies, gathering information about the domain of reality that is to be represented, and determining what kinds of entities are relevant to a representation of that domain given the purpose that this representation is being designed to fulfill.

Determining what the universals and relations amongst universals dealt with in the subject-matter are has to do both with analyzing the subject-matter itself, and with locating the various entities with which it deals in the context of a more general ontology. Concretizing the information that has been gathered in the form of a representational artifact involves providing a

[27] Compare with the very helpful discussion in Eric Little, "A Proposed Methodology for the Development of Application-Based Formal Ontologies", at http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS//Vol-94/ki03rao_little.pdf.

systematic statement, whether in the form of diagrams, a computer file or a written document of some sort, of the entities, relations, terminology etc. that are to be included in the ontology.

Once this has been done, the process of regimentation comes in two steps. The first is to check the universals, relations and domain specific terminology contained in the ontology for logical, philosophical and empirical adequacy, including consistency and human intelligibility. This task is essentially a semantic one which involves making sure that the ontology is maximally effective in representing the universals and relationships that exist in the domain that it is intended to be about. This process should include a good deal of fact-checking, including extensive consultation with domain experts (working biologists, medical practitioners, etc.).

The second step of regimenting the ontology is essentially a syntactical one, and involves translating the terminology, relations, etc. contained in the ontology into a format that is computer tractable, whether this be some fragment of first-order logic, a description logic or something else. Finally, then, comes the process of implementing the formalized representational artifact in some actually functioning computing context.

While the process outlined here is essentially top down in nature, in practice there will be a great deal of interaction and feedback between the different steps of ontology design and construction, with constraints at the computational level and the factual information of the specific scientific domain(s) that are to be modeled influencing decisions that are made about which and how many ontological categories and relations are to be used, and vice versa. The following sections will discuss the steps of this process in greater detail. This will include introducing necessary theoretical notions and clarifications.

Some general principles to be followed in ontology design

Some general principles for ontology construction that underpin the methodology in what follows include:[28]

Realism

'Realism' can be defined as a philosophical position according to which "reality and its constituents exist independently of our (linguistic, conceptual, theoretical, cultural)

[28] For further discussion of these principles, see Pierre Grenon and Barry Smith. "SNAP and SPAN: Towards Dynamic Spatial Ontology." *Spatial Cognition and Computation*, Pierre Grenon, Barry Smith and Louis Goldberg. "Biodynamic Ontology: Applying BFO in the Biomedical Domain." In D. M. PIsanelli (ed.) *Ontologies in Medicine.* Amsterdam: IOS Press, 2004, 20—38, and Barry Smith and Pierre Grenon. "The Cornucopia of Formal Ontological relations." *Dialectica* Vol. 58, 3, (2004), pp. 279—296.

representations thereof."[29] Just as science is the attempt to come to know the general features of reality in the form of universals and relationships obtaining amongst them, so realism is the thesis that the things that scientific knowledge is about are in fact real, mind-independent things. This position has a number of general consequences for ontology. These include the position that ontologies are representations of reality, not representations of people's concepts or mental representations of reality; that science discovers truths about reality, and that the facts that science discovers, such as the nature of universals in the biological or chemical domains, existed long before it occurred to human beings to conceptualize and search for them; and that not all representational schemes are equally good, some are better than others precisely in that they are better representations of reality. There are a number of other principles that go hand in hand with the commitments of realist ontology.[30]

Perspectivalism

Perspectivalism involves the recognition that reality is a complex and variegated phenomenon. While not all representations are good, because some are accurate to the facts of reality and some are not, there are nevertheless many different representations that are equally good (good in the sense of being true) precisely in that they capture different and important features of one and the same reality. These include viewing reality in terms of substances or continuants and their qualities on the one hand, and in terms of occurrents or processes on the other (a distinction that will be more fully developed below), and also viewing reality at different levels of granularity, ranging from the microscopic world of atoms, chemical reactions and molecules, to the macroscopic world of organisms, ecosystems and galaxies.[31]

---

[29] Pierre Grenon and Barry Smith. "SNAP and SPAN: Towards Dynamic Spatial Ontology." *Spatial Cognition and Computation,* p. 138.

[30] For more on realism, see Smith Barry: "From Concepts to Clinical Reality: An Essay on the Benchmarking of Biomedical Terminologies", in *Journal of Biomedical Informatics*, forthcoming; Barry Smith. "Beyond Concepts: Ontology as Reality Representation." In Achille Varzi and Laure Vieu (eds.), *Proceedings of FOIS 2004. International Conference on Formal Ontology and Information Systems,* Turin, 4—6 November 2004; Gunnar O. Klein & 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." http://ontology.buffalo.edu/concepts/ConceptsandOntologies.pdf; Ingvar Johansson. "Bioinformatics and biological reality." *Journal of Biomedical Informatics* 39 (2006) 274—287.

[31] One noteworthy example of perspectivalism in the history of ontology is Edmund Husserl, who held that there were three separate and irreducible material ontological domains, corresponding to the "psychological", the "social" and the "natural" respectively. Husserl maintained that each of these domains had its own governing essence or universal, and its own characteristic internal laws and relationships, none of which could be reduced or explained away in terms of the kinds of things in the other two domains. Nevertheless, Husserl maintained that there were certain formal ontological relationships and categories that applied to entities in all of these domains, and that universals and relationships from each of these domains were coinstantiated in particular objects in the actual

Fallibilism

Fallibilism involves commitment to the idea that, though our current scientific theories are the best candidates we have for the truth about reality, it may nevertheless be the case, given the (realist) fact that reality exists independently of our ways of conceptualizing it, that portions of our current knowledge are incorrect. The fallibilist maintains that it is a matter of empirical investigation what the facts of reality are, and recognizes that empirical investigation is an on-going, open-ended, experimental process.

Adequatism

The approach being developed here also presupposes the principle of adequetism. This is the position that a good theory of reality must do justice to all of the different phenomena that reality contains. In opposition to the tendency to attempt to reductively explain higher level macroscopic phenomena in terms of "more basic" or fundamental components of reality, adequatism entails that the entities in any given domain of reality be taken seriously on their own terms first. Thus, just as an ontology of physics should be about atoms and sub-atomic particles, and an ontology of chemical reactions should include the existence of various kinds of elements and compounds, so an ontology of biological phenomena should include the existence of, at various levels, cells, organs, biological systems and organisms, as well as populations and environments. The goal of adequatism is to do justice to the vast array of different kinds of entities that exist in the world, in different domains and at different levels of granularity, rather than ignoring them or attempting to explain them away.

The following provides a very general over view of the ontology design process. More will be said about many of the headings below in subsequent sections.

A schematic overview of the ontology design process

- *Explicitly determine and demarcate the domain and extent of the ontology.*

1. Formal & material ontologies
   i. Formal Categories: what formal ontological categories are important for the domain?
   ii. Formal Relations: what formal ontological relations are important for the domain?
2. Granularity: what is the appropriate maximum and minimum level of granularity or complexity that the ontology requires?

spatial-temporal world. For some discussion of this, see Smith, D. W. "Mind and Body". Printed in Smith, B. & Smith, D. W. eds. *The Cambridge Companion to Husserl*. New York: Cambridge University Press, 1995.

3. Relevance: what is relevant to include in the ontology given i) the domain the ontology is intended to represent and ii) the purpose for which the ontology is to be used?

- *Determine what the universals and relations amongst them are.*

4. Gathering information: what are the important general terms and relations dealt with in the domain? Organize these into an initial list, preferably with some definition and organization into tentative categories.

5. Scientific investigation: How do the terms and relations function in scientific theories of the domain being represented? Are the terms and relations that have been collected an adequate reflection of what is most crucial for understanding the truth about this domain of reality as reflected in current scientific knowledge?

6. Thought experiments and imaginative variation: What are the essential or defining features i) of the domain as a whole, ii) of the particular entities and relations that have been selected as crucial?

- *Concretize the cognitive representations of these things as representational artifacts (bearing in mind that it is the things, not the cognitive representations of them, that are to be represented by these artifacts).*
- *Regiment the representational artifact for-*

7. Logical, philosophical and scientific coherence.

Terminology.

Definition of terminology.

Taxonomy, especially is_a hierarchies.

Categories and relations, align the domain information with relevant formal ontological categories and relations.

8. Coherence and compatibility with other relevant ontologies.

9. Human understandability.

- Formalize the regimented representational artifact in a computer tractable format.

- Implement the formalized regimented representational artifact.

Chapter 4: Explicitly determine and demarcate the domain and extent of the ontology

The first step in constructing an ontology is to explicitly determine the intended domain of that ontology, to answer the question "what part of reality is this ontology an ontology of?" Providing an explicit statement of the intended subject-matter of an ontology at the outset helps to focus the effort of constructing the ontology by indicating what principles and information need to be included while at the same time ruling out other information as un-important for constructing an ontology of the given domain. For example, the documentation for the Foundational Model of Anatomy, an ontology of human anatomy, reads "The FMA…is strictly constrained to "pure" anatomy, i.e., the structural organization of the body".[32] This statement makes it relatively clear what information is, and what information is not, properly a candidate for inclusion in the FMA, and thus also what terms, ontological categories, universals and relationships might need to be included.

Here we begin by outlining a number of issues that are pertinent to determining the domain of an ontology, including the distinction between formal and material ontologies, and the issues of subject-matter granularity and subject-matter relevance.

Formal and material ontologies

An important distinction to bear in mind when constructing an ontology, and one that will play a role in helping to demarcate the subject-matter of the ontology, is that between formal and material (or domain-specific) ontologies.[33]

A formal ontology is a representation of the categories of objects and of the relationships within and amongst categories that are to be found in any domain of reality whatsoever. The most obvious formal ontological category is "entity"; no matter what science one is considering that science studies entities, and thus the category "entity" applies to the subject matter of that science. A relatively uncontroversial formal ontological relation is the relation "part_of". Even sciences committed to the study of very very small entities acknowledge the existence of at least some parts of those entities, and therefore the relationship "part_of" applies to those domains as well.

A material or domain ontology, by contrast, consists of a representation of the material categories, universals and relationships amongst universals that are to be found in some specific domain of reality, such as genetics, anatomy, plant-biology, cell-biology, physiology, etc. Characteristic of a material ontology is that it will contain many categories, universals and

[32] FMA, http://sig.biostr.washington.edu/projects/fm/AboutFM.html, July 17, 2006.
[33] One of the earliest explicit recognitions of the distinction between formal and material ontologies is to be found in Husserl, E. *Logical investigations Vol. II*. Trans. J. N. Findlay. Amherst: Humanity Books, 1900-01/2000, especially the third and fourth Investigations. The following discussion is also very sympathetic to ideas presented by Alan Rector et. al. "Simple Bio Upper Ontology", http://www.cs.man.ac.uk/~rector/ontologies/simple-top-bio/, especially section "II General Considerations".

relations amongst universals that are not to be found in formal ontologies or in other material ontologies (consider the categories, universals and relationships amongst them that must be dealt with in psychology, physics and sociology respectively).

When one sets out to construct an ontology, it will most often be a material or domain-specific ontology that one is interested in constructing. However, for purposes of managing the sea of biomedical information, the relevance of formal ontology to the construction of material ontologies is two-fold.

First, to the extent that the categories and relations in a formal ontology are well-structured and systematically defined, organizing the universals and relations of a given domain ontology in terms of them will render that domain ontology equally well-structured and hold out at least the possibility that the other universals and relationships in that domain ontology can be well-defined as well. What it means to "organize" a domain ontology in terms of a formal ontology will be discussed in more detail below. However, the basic idea of what this means can be expressed by saying that if, in a formal ontology, the relationship "part_of" has a semantic interpretation such that "if x part_of y and y part_of z, then x part_of z", then in a domain ontology focusing on, for example, anatomy, where a finger part_of arm and arm part_of body, it should be possible to infer that finger part_of body based on the definitions and information explicitly contained in that ontology, and likewise in all similar cases. Thus organizing a material ontology in terms of a formal ontology involves encoding information in the material ontology in a univocal and logically/mathematically precise way based on the fixed meanings of categories and relations in the formal ontology.

Second, using formal ontologies to structure domain ontologies helps to ensure interoperability or communication between and amongst domain ontologies. If the basic set of categories and relations that have been used to structure two different domain ontologies have the same meaning in both ontologies, this will make it easier to bring the information in these ontologies together, and to compare them in various ways. Indeed, one of the primary reasons for interest in formal or "top-level" ontologies in the information science community is the promise that such ontologies hold out of making wide-spread interoperability possible.[34]

---

[34] For an example of the use of formal ontological principles to improve a domain application, see Simon Jonathan, Fielding James Matthew, Dos Santos Mariana Casella, Smith Barry: "Reference Ontologies for Biomedical Ontology Integration and Natural Language Processing", in *International Journal of Medical Informatics*, forthcoming.

For these reasons it is important, at the outset of designing a domain specific ontology, to consider what formal ontological categories and relations might apply to the domain at hand, and to select a formal ontology with sufficient and sufficiently clear categories and relations to handle the basic kinds of entities to be found in the domain in question. It is important to note that, by definition, a formal ontology should not contain all relations and universals that are pertinent to a given domain. Rather, the specific content of a given domain will be "added on" to or "structured in terms of" a formal ontology based on the best available scientific knowledge about that domain.

There are currently a number of different and competing formal ontologies that have been proposed, consisting of various degrees of expressive power, completeness, and implementability.[35] Aside from these recent developments, the history of philosophy itself contains numerous proposals and much argumentation regarding the basic categories and relations that are to be found in reality.[36] Here we will refrain from exploring all of these positions and the various arguments for and against them. Rather, we will begin by briefly discussing some basic features that are common to all formal ontologies, and then by introducing, explaining and developing the Basic Formal Ontology (BFO), which has been developed at the Institute for Ontology and Medical Information Science in Saarbrücken, Germany. This latter example will then be used as we develop some of the more specific principles for good ontology design and best ontology practices. Importantly, the majority of best ontology practices to be discussed in what follows do not depend in any way on the

---

[35] Claudio Masolo et al., *Wonderweb Deliverable D18* (2003 [cited May 6, 2006]); available from http://wonderweb.semanticweb.org/deliverables/documents/D18.pdf, see also Gangemi, A. et. al. "Sweetening Ontologies with DOLCE", http://citeseer.ist.psu.edu/cache/papers/cs/26864/http:zSzzSzwww.ladseb.pd.cnr.itzSzinforzSzOntologyzSzPaperszSzDOLCE-EKAW.pdf/gangemi02sweetening.pdf, and Heller, B. & Herre, H. "Ontological Categories in GOL", *Axiomathes* 14: 57—76, 2004, http://www.onto-med.de/en/theories/gfo/index.html; Alan Rector et. al. "Simple Bio Upper Ontology", http://www.cs.man.ac.uk/~rector/ontologies/simple-top-bio/.

[36] Aristotle. *Metaphysics*. In McKeon, R. ed.*The Basic Works of Aristotle*. New York: The Modern Library, 2001, Nelson Goodman. *Problems and Projects*. New York: Bobbs-Merrill Company, 1972, Husserl, E. *Logical Investigations Vol. II.* Trans. J. N. Findlay. Amherst: Humanity Books, 1900-01/2000, Kant, I. *Critique of pure reason*. Trans. Paul Guyer and Allen W. Wood. Cambridge: Cambridge University Press, 1871/1998, Lowe, E. J. *A Survey of Metaphysics*. Oxford: Oxford University Press, 2002; Mertz, D. W. *Moderate Realism and its Logic*. London: Yale University Press, 1996, Quine, W. V. O. *Word and object*. Cambridge: MIT Press, 1960, Simons, P. *Parts: A Study in Ontology*. Clarendon: Oxford University Press, 1987, Smith, B. ed. *Parts and Moments*. Munchen: Philosophia Verlag, 1982, Varzi, A. 'Mereology', in Edward N. Zalta (ed.), *Stanford Encyclopedia of Philosophy*, Stanford: CSLI (internet publication), 2003. ftp.

categories and relations of BFO in particular, but are applicable to the construction of any domain ontology, based on almost any formal ontology whatsoever.

Basic features of formal ontologies

Formal ontologies specify reality along two essential dimensions. First, they state what the basic categories of reality are. Second, they state the basic relations that hold within and among objects belonging to the basic categories of reality.

The term 'ontological category', as it is being used here can be interpreted to mean "universal that applies to every material domain of reality," in the case of formal ontologies, and "universal of highest generality within a domain" in the case of domain ontologies.[37]

The ontological categories in a formal ontology should include all and only very basic and very general kinds of entities. Examples of ontological categories include the category "universal" and the category "particular" that have already been discussed. Other ontological categories might include "individual", "collective", "property", "dependent entity", "process", "event" or "spatial region". Each of the categories in an ontology should be carefully defined, ideally involving the statement of identity conditions of the form "an x belongs to category y just in case it is a, b, and c…", and no two categories should be identical or have the classes that they designate overlap (no entity in the world should belong simultaneously to two different formal ontological categories).

A good test of whether or not something is a formal ontological category is to try and find two universals from very different scientific domains, both of which can be correctly understood as belonging to the category in question. Thus the fact that cups, rocks, planets, organisms and perhaps even some social organizations can be described as substances, meaning as independently existing entities that maintain their identity through the loss and replacement of (most kinds of) parts and are also the bearers of qualities or properties of various sorts, makes "substance" a likely candidate for a formal ontological category.

---

[37] The use of the term 'category' has a long and variegated usage in the history of philosophy, and new uses are being developed daily, both by philosophers and by researchers in the information sciences and in bioinformatics. We will not discuss all of the different possible uses of the term 'category' here. For an introduction to the philosophical usage and issues surrounding the term, see Thomasson, A., "Categories", *The Stanford Encyclopedia of Philosophy* (Winter 2005 Edition), Edward N. Zalta (ed.), URL = http://plato.stanford.edu/entries/categories/.

The formal relations contained in an ontology are extremely important. They serve to specify both the internal structure of objects that are to be found within ontological categories and the relationships that obtain between different ontological categories.

An example of a relation performing the first kind of function is the "part_of" relation. For any given object falling within any given ontological category, it is possible to further analyze that object by analyzing its parts.

An example of a relation performing the second kind of function would be the relation "instantiates". It is entities belonging to the ontological category "particular" that "instantiate" entities belonging to the ontological category "universal". Thus "Fido instantiates Golden Retriever", and "this-oxygen-molecule instantiates oxygen". An ontology that includes the categories "particular" and "universal", as well as the relationship of "instantiates" holding between them is both more explicit and more structured than one that only contains these categories. Further, it should be noted that the relationship "instantiates" can be used to help define each of these categories. For example, one could begin to define 'universal' as "an entity that can be instantiated", and one can begin defining 'particular' as "an entity that instantiates universals, but is not instantiated by any other thing." In such cases adding explicit formal relations to a system of ontological categories helps to implicitly define the nature of these categories themselves.

As with formal categories, so with formal relations, a good test of whether or not a relation is formal is to check and see if it applies to entities in many different scientific domains. The fact that cells, buildings, molecules, organisms, environments and football games can all be said to have "parts" in the same sense makes the "part_of" relationship a good candidate for inclusion in a formal ontology.

The dimensions of an ontology

Whereas formal categories and relations are intended to be basic features of any domain ontology whatsoever, the constraints imposed by a formal ontology alone do not yet help to determine what the specific content (in terms of domain specific universals and relations) to be included in a domain ontology should be. The question of what information should, and what information should not, be included in a given domain ontology is left unanswered by formal ontological considerations alone. There are at least two general dimensions along which the appropriate contents of a domain specific ontology need to be determined. These two dimensions might be called the "horizontal" and the "vertical" dimensions of a domain ontology respectively, and we will discuss them here under the headings of "relevance" and "granularity".

Relevance

The "horizontal" demarcation of the contents of an ontology is the problem of determining what and how much existing information about a given domain should be included in that ontology, and can be summarized by describing it as a problem of determining what is *relevant* for the ontology. An ontology should include only that information about reality that is relevant to the domain that it is a representation of.[38] The issue of relevance is a complicated one about which much has been written. Though there is no clear consensual definition of 'relevance' applying to all cases of its use, this does not mean that what is relevant to a given ontology is a merely subjective affair, to be decided purely by the attitudes, purposes or opinions of the ontology designers. Rather it is possible to understand relevance as an objective phenomenon, one tied to basic features of reality and of the various sciences that study those features of reality. In this connection there are two considerations that are especially important to keep in mind when determining what is relevant to a given ontology.

The first is that what is relevant to the ontological representation of a given domain will be determined by reality itself. Domains of reality contain their own internal objective relevancies. What is objectively relevant to a gene ontology, for example, will be determined first and foremost by the nature of genes themselves, what they are, what processes they characteristically initiate or are involved in, etc.

Similarly, an ontology explicitly intended to include the scientific information that we have about cells would be expected to include at least a complete listing of known types of cells (eukaryotic, prokaryotic, etc.) and a complete specification of the characteristic kinds of parts of each of these kinds of cells. However, while it is relevant to such an ontology to include the information that genetic material is located in cells, it is arguably not relevant to such an ontology to include detailed information about the nature of genes, their characteristic expressions or the processes involved in their recombination and transmission; this information would be much more relevant to an ontology devoted specifically to the representation of the

[38] Ceusters, Werner & Smith, Barry. "A realism Based Approach to the Evolution of Biomedical Ontologies." in *Proceedings of AMIA Symposium 2006,* forthcoming.

nature of genes and genetic processes themselves. Importantly, these relevancies are determined by the nature of genes and cells themselves: understanding the kinds of cells there are and the parts and processes that they are involved in, in short understanding the nature of cells, does not require understanding everything that there is to know about genetics. Thus the vast majority of information that is possessed by the science of genetics is not directly relevant to the representation of the ontology of cells, and so should not be included.

The fact that the nature of cells is different from the nature of genes suggests that there should be (at least) two separate domain ontologies, one for cells and one for genes. The way to handle the fact that there are connections between cells and genes is simply to build explicit links and references from the one ontology to the other in those places where the information contained in one comes to an end, and the information in another begins. Assuming that these domain ontologies have been structured in terms of the same formal ontology, and that terminological usage in the two is consistent, it should be possible to connect them and to bring them into alignment in this way.

The second major consideration to be kept in mind regarding relevance is that what is relevant to an ontology will also be determined by the purpose for which that ontology is being designed. Goal-oriented human activities generate their own sets of objectively relevant features of reality, actions, etc. For example, the project of going to the grocery store to buy milk automatically makes the following entities relevant: the store itself, the means of transportation, the transportational route, the milk, some means of payment, and an ordering amongst the series of activities that must be engaged in. The same can be said for scientific investigations (the goal of which is to gain new information about the world) and medical treatments (the goal of which is to heal a patient based on existing biological knowledge and medical techniques), though of course the number of things relevant to these goal oriented activities will be greater and involve more complex relations.

<u>Reference and application ontologies</u>

Recognizing these two kinds of relevancy substantiates the distinction between a reference ontology and an application ontology. A reference ontology is a representational artifact analogous to a scientific theory, in which expressive completeness and adequacy to the facts of reality are of primary importance. An application ontology, on the other hand, is a representational artifact designed to assist in the achievement of some specific goal. Reference ontologies will be constructed and structured primarily based on objective relevance, while application ontologies will be constructed and structured primarily in terms of goal-oriented relevancies. Ideally, application ontologies will make use of (and also be able to reuse) portions of reference ontologies in order to accomplish their particular ends.

When designing an ontology then, it is essential at the very beginning to carefully consider what the domain of the ontology is. This will include asking the question of whether the ontology is primarily intended to be a comprehensive representation of scientific information in a given domain, and thus a reference ontology, or whether it is primarily intended to be used by human beings in order to accomplish certain very specific goals, such as medical treatments, in which case it is an application ontology. In the former case what is relevant to include in the ontology will depend primarily on the nature of the objects that exist in the domain to be represented, while in the latter case what is relevant to the ontology will depend on what portions of reality and what information is objectively salient to achieving the stated objective or goal.[39]

Granularity

In a sense, the issue of determining the appropriate granularity of an ontology is just a sub-part of the problem of determining what is relevant to be represented in that ontology. However, whereas we have described the problem of relevance as the horizontal problem of determining how much and what kind of information to included in an ontology, the problem of granularity is the specifically vertical problem of determining how fine-grained or course-grained the ontology should be. The issue of granularity arises because things in reality, and also the parts of things, come in many different sizes and possess varying degrees of complexity. There is a continuum spanning from the level of sub-atomic particles, atoms and molecules, through the levels of ordinary everyday objects such as animals, rocks, tables and lakes, and on to the level of ecosystems, planets, solar systems, galaxies, and ultimately the universe itself. Things exist at all of these different levels, and as things become bigger and more complex (consisting of ever more and more diverse parts) they also exhibit features not to be found in

[39] See Christopher Menzel. "Reference Ontologies / Application Ontologies: Either / Or or Both/And?" in P. Grenon, C. Menzel, and B. Smith (eds.), *Proceedings of the Workshop on Reference Ontologies and Application Ontologies*, KI 2003, September 16, 2003, Hamburg, Germany. CEUR Workshop Proceedings, vol. 94 (2004), ISSN 1613-0073, & Fielding, Matthew James, Mariana Casella Dos Santos, and Barry Smith: "Reference Ontologies for Biomedical Ontology Integration and Natural Language Processing", in *International Journal of Medical Informatics*, forthcoming.

smaller less complex entities.[40] The problem of granularity for ontology design is the problem of deciding the upper and lower limits of entity size and complexity that are to be represented in a given domain-ontology. Should an ontology of mountains include information about the atoms out of which the mountains are composed, should an ontology of the circulatory system include detailed information about blood cells, and should an ontology of cells include detailed information about genetics? Each of these questions is a question about the appropriate granularity of the respective ontologies mentioned. Importantly, any given ontology should specify both an upper ('largest') and a lower ('smallest') boundary or level of granularity that it will represent.[41]

The appropriate granularity for an ontology should be decided, essentially, along the same lines as should be used in deciding what is relevant for inclusion in the ontology in general. Namely, in the case of reference ontologies by focusing on the nature of the objects to be represented, and in the case of application ontologies by focusing on the objective saliencies determined by the intended goal. There is, however, an important confusion that needs to be avoided when determining the granularity of an ontology.

Determining the granularity of an ontology (the size and complexity of the kind of objects that it represents) is not the same as determining the generality of an ontology (the abstractness of the universals that it represents).

Suppose that there was good reason to develop an ontology of human anatomy that did not represent anything beneath the level of granularity possessed by bodily organs such as the heart, kidneys, etc. In such a case, the ontology would represent universals and relations amongst them for all features of the body down to the level of organs, but nothing further. Now, this ontology would also include a universal representation for "human heart", and importantly "human heart" is (arguably) the most specific (least general) universal that can be discussed before one crosses the universal/particular divide and must begin discussing the existence of particular hearts instantiating this universal. Such an ontology would then be very "specific", in

---

[40] It is worth noting here that there is no guaranteed relationship between the size of an object and its complexity. Genes are very small and yet very complex, whereas mountains are very big, but relatively simple in their composition and functioning. Nevertheless, the issue of granularity encompasses both the literal question of what sizes of objects should be represented in an ontology and also what level of complexity should be represented, and thus the two issues are mentioned together here.

[41] For more on granularity and the theory of granularity, see Kumar Anand, Smith Barry, Novotny Daniel: "Biomedical Informatics and Granularity", in: Functional and Comparative Genomics, 5 (2004) , 501-508, and the papers at the following link http://ontology.buffalo.edu/smith/articles/vagueness.html.

that it would represent universals falling at the bottom of their respective hierarchies, but it would also have a relatively coarse granularity (insofar as there is more to human anatomy than just the level of the whole body through the level of bodily organs).

Alternatively, there might be reason to develop an ontology of genes that contained only information common to all genes, no matter what particular living thing they are found in. Such an ontology might well leave out of account many features of genes that exist only in particular species of organisms, but not in all living things. This ontology would be very general in the sense that the universals it represented would be higher in the taxonomic hierarchy of gene-universals, but it would (or at least could) still have a very fine grained granularity, insofar as it is representing things that exist at the level of genes, their components, their properties and the processes that they engage in.

From this it should be clear that the task of determining the level of granularity that an ontology will represent (how fine- or course-grained it will be) is separate from the task of determining the level of generality that that ontology will represent (in terms of the abstractness of the universals that it contains).[42]

Chapter 5: Introduction to Basic Formal Ontology and Relations

Fully structuring domain ontologies in a way that makes them computer-tractable and interoperable, as well as in a way that renders the information that they contain as clear, rigorous, and unambiguous as possible, requires the use of formal or "top-level" ontologies. In addition, the important issues of terminology selection, term-definition, and classification in the sense of the construction of tree-like structures organizing the information contained in a domain ontology all benefit from and can be better understood in the context of, an explicitly defined formal ontology. Here we provide an introduction to some of the core features of the Basic

[42] For a conflation of these two issues, see Natalya F. Noy and Deborah L. McGuinness. ``Ontology Development 101: A Guide to Creating Your First Ontology''. Stanford Knowledge Systems Laboratory Technical Report KSL-01-05 and Stanford Medical Informatics Technical Report SMI-2001-0880, March 2001. "Deciding where classes end and individual instances begin starts with deciding what is the lowest level of granularity in the representation. The level of granularity is in turn determined by a potential application of the ontology. In other words, what are the most specific items that are going to be represented in the knowledge base?" http://ksl.stanford.edu/people/dlm/papers/ontology-tutorial-noy-mcguinness.pdf