Chapter 1
The Interplay Between Ontology as Categorial Analysis and Ontology as Technology

Roberto Poli and Leo Obrst

1.1 Introduction

The notion of ontology today comes with two perspectives: one traditionally from philosophy and one more recently from computer science. The philosophical perspective of ontology focuses on categorial analysis, i.e., what are the entities of the world and what are the categories of entities? *Prima facie*, the intention of categorial analysis is to inventory reality. The computer science perspective of ontology, i.e., ontology as technology, focuses on those same questions but the intention is distinct: to create engineering models of reality, artifacts which can be used by software, and perhaps directly interpreted and reasoned over by special software called inference engines, to imbue software with human level semantics. Philosophical ontology arguably begins with the Greek philosophers, more than 2,400 years ago. Computational ontology (sometimes called "ontological" or "ontology" engineering) began about 15 years ago.

In this chapter, we will focus on the interaction between ontology as categorial analysis ("ontology_c", sometimes called "Big O" ontology) and ontology as technology ("ontology_t", sometimes called "Little o" ontology). The individual perspectives have each much to offer the other. But their interplay is even more interesting.[1]

This chapter is structured in the following way. Primarily we discuss ontology_c and ontology_t, introducting notions of both as part of the discussion about their interplay. We don't think they are radically distinct and so do not want to radically distinguish them, intending by the discussion of the interplay to highlight their distinctions where they occur, but thereby emphasize their correspondences, and, in fact, their correlations, complementarities, interdependencies. They are distinct

R. Poli (✉)
University of Trento, Trento, Italy
e-mail: Roberto.Poli@unitn.it

[1] Cf. Daconta et al. (2003, p. 186). The first use of this "Big O, little o" terminology, as known by the authors, is in Guarino (1995). The distinction made between ontology_c and ontology_t is first made in Poli (2001b).

R. Poli et al. (eds.), *Theory and Applications of Ontology: Computer Applications*,
DOI 10.1007/978-90-481-8847-5_1,

perspectives after all, we want to emphasize, not distinct analytical methodologies, nor do they provide distinct analytical products. We discuss some of the historical definitions of ontology_t, as they emerged during the 1990s. We then provide our own take on the nature of ontology_t. As part of this exposition, we briefly discuss the levels and representation of ontologies, ranging over the typical levels of upper ontologies, middle ontologies, and domain (and sub-domain) ontologies.

Although we cannot discuss the knowledge representation languages typically used by ontology_t, from Semantic Web languages such as OWL[2] (primarily a description logic) to First-Order Logic (FOL, predicate calculus) languages such as ISO Common Logic,[3] nor the automated reasoning over ontologies that is of potential benefit to ontology_c as well as to ontology_t, we consider these issues important but better exposed in another venue. The interested reader is, therefore, directed to Chapter 2, Ontology Architecture, for a fuller exposition.

We do, however, lay down some principles by which we believe ontologies_t should be developed, based on analysis from ontology_c, and introduce the notion of "levels of reality". We illustrate the interplay of the two notions of ontology by providing an extended discussion of ontological entities in a hypothetical biology ontology.

Finally, we conclude by looking to the increasing interaction between these two aspects of ontology in the future. We briefly discuss some common problems which require the interplay of ontology_c and ontology_t, and which will assume much greater prominence once the more basic issues are elaborated on and scientific concensus established, i.e., ontology modularity, mapping, context-determination and representation, and vagueness and uncertainty. Ontology_t needs to be informed by ontology_c and its analytical methods. Ontology_c will increasingly benefit from the sound and consistent software engineering products arising from ontology_t.

1.2 Ontology_c

Ontology and ontologies have been given many different definitions, both on the philosophical side and on the technological side.

From the perspective of categorial analysis in philosophy, ontology has been viewed as both a part of metaphysics and as a part of science. Historically, ontology has been a branch of metaphysics, interested in formulating answers to the question of what exists, i.e., what's the inventory of reality, and consequently in defining categories (kinds) of entities and the relationships among the categories. Metaphysics asks different questions than does ontology, notably the question about the nature of being as a whole.

In our understanding, ontology should also be viewed as following along the same path as science, i.e., that ontology organizes and classifies the results from that which science discovers about reality. Furthermore, ontology not only depends

[2] Bechhofer et al. (2004).

[3] ISO Common Logic: Common Logic Standard. http://cl.tamu.edu/.

on science but can also provide tools for the clarification of science itself, in the form of ontologically clarified and reconstructed sciences. Ontology and science can therefore support one another.

A further point of contention or at least confusion is that between ontology and epistemology, i.e., on the study of what is vs. the study of what is ascertained and how it is ascertained. Ontology requires knowledge about what is, and if knowledge is described as, for example, justified belief, then ontology may be thought to devolve to knowledge and from thence to belief and justification for belief, i.e., the realm of evidence, manners and methods by which one adjudicates evidence to form belief, and thus epistemology.

Ontology is not epistemology, but has a complex relationship to epistemology. Ontology is primarily about the entities, relations, and properties of the world, the categories of things. Epistemology is about the perceived and belief-attributed entities, relations, and properties of the world, i.e., ways of knowing or ascertaining things. So epistemology is about empirical evidence gleaned that will be described or characterized by ontology.

Contemporary ontology can be characterized in a number of ways, all of which can be considered layers of theory (Poli, 2003):

(1) *Descriptive ontology* concerns the collection of *prima facie* information either in some specific domain of analysis or in general.
(2) *Formal ontology* distills, filters, codifies and organizes the results of descriptive ontology (in either its local or global setting). According to this interpretation, formal ontology is formal in the sense used by Husserl in his *Logical Investigations* (Husserl, 2001; originally 1900–1901). Being "formal" in such a sense means dealing with categories like *thing*, *process*, *matter*, *form*, *whole*, *part*, and *number*. These are pure categories that characterize aspects or types of reality and still *have nothing to do with the use of any specific formalism*.
(3) *Formalized ontology:* Formal codification in the strict sense is undertaken at this third level of theory construction. The task here is to find the proper formal codification for the constructs descriptively acquired and categorially purified in the way just indicated. The level of formalized constructions also relates to evaluation of the adequacy (expressive, computational, cognitive) of the various formalisms, and to the problem of their reciprocal translations. In this sense, formalized ontology refers to the actual formalization of ontology in a logical language, typically but not always First Order Logic (FOL). In ontology_t, this could be rendered in a knowledge representation language such as the FOL-based ISO Common Logic or in the description logic-based Web Ontology Language OWL.

The close similarity between the terms "formal" and "formalized" is rather unfortunate. One way to avoid the clash is to use "categorial" instead of "formal".[4]

[4]Note the philosophical, common use of "categorial" instead of the term "categorical" employed in this chapter, which comes closer however to the mathematician and logician's use of the term "categorical", as for example in Category and Topos Theory.

Most contemporary theory recognizes only two levels of work in ontology and often merges the level of the formal categories either with that of descriptive or with that of formalized analysis. As a consequence, the specific relevance of categorial analyses is too often neglected.

The three levels of ontology are different but not separate. In many respects they affect each other. Descriptive findings may bear on formal categories; formalized outcomes may bear on their formal equivalents, etc. To set out the differences and the connections between the various ontological facets precisely is a most delicate but significant task (Poli, 2003).

1.3 Ontology_t

Ontological engineering, i.e., ontology from the perspective of computer science, has issues comparable to that of philosophical ontology, but reflected technologically in the attempt to develop ontologies as software usable models. So ontology from the perspective of computer science is both a computer science and a computational or software engineering problem. On the one hand, "ontological engineering"[5] historically had its origins as an engineering problem, as an attempt to create software usable models of "the ways things are, with the things that are" to endow software with human level representations of "conceptualizations" or semantics. On the other hand, there are efforts that intend to make an "ontological science", as for example, that of the National Center for Ontological Research (NCOR) (Obrst, Hughes and Ray, 2006).[6] Such an effort would include strong evaluation criteria and possibly ontology certification.

Although having antecedents in the late 1980s, as formal ontology in philosophy and formal semantics in linguistics began to impact computer science and especially artificial intelligence, ontological engineering as a discipline can be marked as originating approximately in 1991, with Neches et al. (1991) reporting on the United States Defense Advanced Research Projects Agency's (DARPA) Knowledge Sharing Initiative, and Gruber (1991), followed soon after by work by Gruber (1993), Guarino (1994), and Guarino and Poli (1995).

1.3.1 Ontology_t Definitions

The first proposed definition of ontology in computer science was that of Gruber (1993)[7]: "an ontology is an explicit specification of a conceptualization", which

[5]The first occasion of use of the term "ontological engineering" is apocryphal: perhaps it occurred as part of the Cyc project (Guha and Lenat, 1990).

[6]National Center for Ontological Research (NCOR): http://ncor.buffalo.edu/.

[7]Anecdotally, the term "ontology" had been used in computer science and artificial intelligence since the late 1980s. One of the authors of this chapter described the use of ontologies and rules in Obrst (1989).

was intended to contrast with the usual definition of ontology in philosophy, i.e., to emphasize that what was being talked about was ontology_t in our terminology: ontology as a computational engineering product. The notion of "conceptualization" was defined in Genesereth and Nilsson (1987) to be "the objects, concepts, and other entities that are presumed to exist in some area of interest and the relationships that hold them" (Gruber, 1993) and presumably the "area of interest", now typically called "domain", is a portion of the world.

Guarino and Giaretta (1995) took up the challenge to clarify what was meant by this and other emerging definitions of ontology_t. In Guarino's and Giaretta's analysis, there were a number of ways to characterize ontology (quoted from Guarino and Giaretta, 1995, p. 25):

1. Ontology as a philosophical discipline
2. Ontology as an informal conceptual system
3. Ontology as a formal semantic account
4. Ontology as a specification of a conceptualization
5. Ontology as a representation of a conceptual system via a logical theory

   5.1 characterized by specific formal properties
   5.2 characterized only by its specific purposes

6. Ontology as the vocabulary used by a logical theory
7. Ontology as a (meta-level) specification of a logical theory.

By way of a summary: "ontology: (sense 1) a logical theory which gives an explicit, partial account of a conceptualization; (sense 2) synonym of conceptualization." (Guarino and Giaretta, 1995, p. 32)

Note that characterization (4) invokes Gruber's definition. Part of Guarino's and Giaretta's explication involves analyzing Gruber's [derived from Genesereth and Nilsson's (1987)] notion of a conceptualization as being extensional. Instead, Guarino and Giaretta (1995) argue that it should be an intensional notion. Rather than Genesereth and Nilsson's (1987) view of conceptualization as "a set of extensional relations describing a particular state of affairs," in Guarino's and Giaretta's view, it "is an intensional one, namely something like a conceptual grid which we superimpose on various possible states of affairs." (Guarino and Giaretta, 1995) The definition that Guarino and Giaretta end up with is that an ontology is an ontological theory, and as such that it "differs from an arbitrary logical theory (or knowledge base) by its semantics, since all its axioms must be true in every possible world of the underlying conceptualization."

1.3.2 Ontology_t and Epistemology

A further issue about ontology and epistemology should be brought out now, as it relates to ontology_t. We have mentioned that epistemology deals with how knowledge is known. How do my perception and understanding, my beliefs, constrain my arrival at real knowledge or assumed belief, i.e., evidence, knowledge hypotheses prior to their becoming theorems about knowledge (and there should be a clear path from hypothesis to theorem to true theorem, but often there is not). So if an ontology is a theory about the world, epistemology addresses the ways of acquiring enough knowledge (and the nature of that) so that one can eventually frame a theory. In ontology_t, the engineering artifact of the ontology model (a theory) will require epistemological linkage to data. That data can be inaccurate, contain uncertainties, and lead to partially duplicate but inconsistent instances of ontology classes. Epistemology thus is employed in the use and qualification of data and as stored in databases or tagged or indexed in documents.

If ontology states that human beings have exactly one birth date, the data about a specific person is epistemological: in a given set of databases the person instance named John Smith (we assume we can uniquely characterize this instance) may have two or more attributed birth-dates, not one of which are known to be true. Epistemological concerns distort and push off needed ontological distinctions. Evidence, belief, and actual adjudication of true data is epistemological. What the real objects, relations, and rules are of reality are ontological. Without ontology, there is no firm basis for epistemology. Analysts of information often believe that all is hypothesis and argumentation. They really don't understand the ontological part, i.e., that their knowledge is really based on firm stuff: a human being only has one birth date and one death date, though the evidence for that is multivarious, uncertain, and needs to be hypothesized about like the empirical, epistemological notion it is.

In fact, much of so-called "dynamic knowledge" is not ontological in nature (ontological is relatively static knowledge), but epistemological. What is an instance that can be described by the ontology? How do I acquire and adjudicate knowledge/evidence that will enable me to place what I know into the ontological theory? Instances and their actual properties and property values at any given time are dynamic and ephemeral (this particular event of speaking, *speaking_event_10034560067800043,* just occurred; however the speaking_event ontology class has not changed).

1.3.3 Ontology_t as Theory with Philosophical Stances

Ontology_t often considers an ontology to be a logical theory about some portion of the world.[8] Philosophical stance towards theories is therefore quite important, because a given ontological engineer will typically imbue his ontology_t engineering model with constructs aligned with his or her philosophical stance, e.g. as to the preferred theory of universals (nominalism, conceptualism or realism).

---

[8]See for example, the discussion of what an ontology is on the Ontolog Forum site: http://ontolog.cim3.net/cgi-bin/wiki.pl?, i.e., Obrst (2006).

1.4 Interplay Between Ontology_c and Ontology_t

The issues we discuss in this section involve the complex interplay between ontology_c and ontology_t. Two main points are discussed: (1) the proper way of developing formalized ontologies; (2) an illustration of one case in which the interplay between philosophy and computer science can be explicitly seen. We discuss the problem of the "natural" boundaries of a domain ontology and how different types of domain ontologies should be distinguished.[9]

We take both Guarino's and Giaretta's position (Guarino and Giaretta, 1995; ontology as interpreted formal system, i.e., a logical theory) and Gruber's position (Gruber, 1993; ontology as specification of a conceptualization) as problematic. Concerning the former, we think focusing on ontology as interpretation only is insufficient. As reflected in Guarino and Giaretta (1995), in this view ontology is more focused on the interpretation (semantics) of a logical theory, i.e., has more of a conceptual-flavored and model-theoretic position ultimately. A consistent logical theory can be developed about nonsense, for example, with no intent to describe a portion of the real world, the task of philosophical ontology as we see it. Subsequent discussions by Guarino (e.g., Guarino, 1998a, 2002; Masolo et al., 2003) have pointed to a better reconciliation between logical theory and realist-based formal ontology, which is more closely aligned with our view, as discussed next.

Against both of these views, however, we would rather say that ontology starts to be something relevant only when specifically ontological axioms are added to some formal basis (say, FOL). The definition of new concepts without the introduction of new axioms has limited value. In this regard, we consider as exemplar the General Formal Ontology (GFO) (Herre et al., 2006).[10] But we also admire other upper or foundational ontology efforts which have sought to axiomatize their distinctions, including, Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE),[11] Basic Formal Ontology (BFO),[12] Object-Centered High-level Reference Ontology (OCHRE), Suggested Upper Merged Ontology (SUMO),[13] Upper Cyc,[14] etc. Recently there was an effort to reconcile or at least map among many of these upper or foundational ontologies by the Ontolog Forum (Obrst et al., 2006).[15]

---

[9] We do not discuss ontological layers here in any detail. The interested reader instead is pointed toward the chapters on the Categorial Stance and on Ontological Architectures in this volume.

[10] General Formal Ontology (GFO): http://www.onto-med.de/en/theories/gfo/index.html. See also Herre's chapter in this volume.

[11] For DOLCE and OCHRE, see Masolo et al. (2003) and the site: http://www.loa-cnr.it/DOLCE.html.

[12] Basic Formal Ontology (BFO): http://www.ifomis.uni-saarland.de/bfo.

[13] Suggested Upper Merged Ontology (SUMO): http://www.ontologyportal.org/.

[14] Upper Cyc: http://www.cyc.com/cycdoc/vocab/vocab-toc.html.

[15] Ontolog Forum: http://ontolog.cim3.net/cgi-bin/wiki.pl?.

1.4.1 Developing Formalized Ontologies

Concerning the second issue, the interplay of ontology_c and ontology_t, we provide an example that illustrates some domain ontology distinctions. We initially assume that the basic distinction between domain ontologies (DO) and upper ontologies (UO) are given. We further assume that a boundary has been established between a selected UO and the DOs which are or could be subsumed under it. Consequently, it should be clear which concepts pertain to the UO and which pertain to its DOs. Typically, highly general concepts like "process", "part", and "boundary" are likely to be included in a UO, while concepts like "gene", "cell" and "membrane" are likely to be included in a domain ontology for, say, biology. Note that we do suppose that a domain ontology for biology may be considered a domain-specific UO, since the constructs of the domain ontology may correctly have to be made general enough to encompass prospectively an entire science. Considering biology as a domain with respect to a true UO, then in turn a biology domain ontology may be considered a domain-specific UO with respect to many complex sub-domains. These sub-domains can be considered domains in their own right (perhaps also incorporating other domain ontologies, say that of public administration for the case of public health), given the complexity of their subject matter, e.g., mammalian anatomy, neuropathology, genetic engineering, clinical medicine, public health, pharmacology, etc. We might call such a domain-specific UO a middle ontology (that spans multiple domains), a "superdomain" ontology, or simply a domain-specific UO.

Figure 1.1[16] depicts the basic layers and the nomenclature we employ. By "utility ontology" in the above, we mean an ontology that represents commonly used

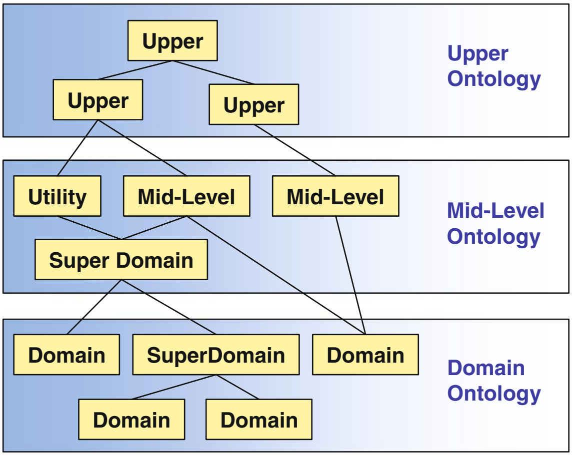

Fig. 1.1 Ontology layers

[16]From Fig. 9.1 in Chapter 8, Ontological Architecture; also see Semy, Pulvermacher and Obrst (2005, p. 8).

concepts, such as Time and Location. However, there is no crucial distinction between a "utility" and a "mid-level" ontology. We do note that in general, a mid-level ontology more concretely represents concepts that are defined more abstractly at the UO level.

To establish our ideas, the two following situations offer useful hints.

Case 1. At the beginning of the previous century the Polish-Russian philosopher of law Leon Petrazycki called attention to a basic theoretical requirement of theory development. We quote one relevant passage:

> "Many theories, comprising no fallacy, are yet inadequate: one may form the concept of "a cigar weighing five ounces", predicate about that class everything known about material things in general (about solid bodies in general, the chemical properties of the ingredients of these cigars, the influence of smoking them on health, and so on); these "theories" while perfectly correct are manifestly inadequate since what is predicated with respect to "cigars weighing five ounces" is also true of innumerable objects which do not belong to that class, such as cigars in general. A theory may be inadequate either (1) because the predicates are related to classes which are too narrow . . . or (2) because the predicate is related to a class which is too broad (such as various sociological theories which attribute "everything" to the influence of one factor which in fact plays a much more modest part)" (Petrazycki, 1955, p. 19).

We may well read "ontologies" where Petrazycki writes "theories". In fact, one may well read "concepts", since cognitive science has a comparable notion concerning "concepts", i.e., that they be non-profligate in a similar manner; especially with respect to what is called the "theory–theory of concepts", concepts as "mental theories" (Laurence and Margolis, 1999, p. 43), and with respect to profligacy: the potential concepts "the piece of paper I left on my desk last night", "frog or lamp", "31st century invention" (Laurence and Margolis, 1999, p. 36). Interestingly, it seems conceptual analysis is recapitulating ontology and semantics, since the former is also addressing categorization, analyticity, reference determination, and the notion of a "prototype" including the notion of "evidential" vs. "constitutive" properties (Laurence and Margolis, 1999, p. 33), which stumbles on the epistemology vs. ontology conundrum.

The point here is that an ontology (viz. a *domain* ontology) may then be inadequate if its boundaries are badly cut. But how should one know where to draw "natural" or "appropriate" boundaries? Some will say that many ontologies at the domain and middle levels correspond to scientific disciplines, i.e., that science and scientific theories apportion the areas of interest. This is partially true, but of course it dismisses intuitive or common-sense ontologies that humans may have, each even considered a logical theory about the world, because they are based on non-scientific generalizations. A theory of parenthood, for example, may not be scientific *yet*, i.e., not based on a combination of sociology, anthropology, biology, psychology, economics, political science, and everything else that might be scientifically known, but it may be a reasonable approximation of reality for the short or mid term, since it's very doubtful those combination of scientific theories will be reconciled anytime soon. This point argues for the inclusion of commonsense theories in lieu of established scientific theories, the latter which may not ever be forthcoming.

Case 2. It is well known that for decades classification theory has labored under an unresolved split between two different methodologies. This split is particularly pronounced in the case of frameworks elaborated by librarians, where it takes the form of the difference between enumerative or taxonomical (Dewey Decimal Classification (DDC), Library of Congress Classification (LCC), etc) classification and faceted classification, also called colon classification, originating from the analysis of Ranganathan (1962). Since faceted classifications are now being proposed for Web construction and other computer-based applications as a more effective way to organize information, a proper understanding of its nature is becoming increasingly relevant. Unfortunately, what is not clear is whether any criteria is available for deciding whether an enumerative or a faceted style of classification should be adopted.

Where lies the *natural* boundary of an ontology_t? The question is both difficult and subtle. May it not be that the boundaries of an ontology may depend on subjective intentions? Just as semantics-pragmatics in linguistics and philosophy of language represents a spectrum, with the latter pole being focused on "semantics in context, with respect to a given use and intent", there are subjective intentional issues here. However, the problem of subjective reasons or needs ("I am developing this ontology for this and that reason, based on these use cases") subtly misses the point. Subjective motivations are always there, and they may more or less severely constrain the ontology to be build. We think that even moderate subjectivism is problematic. So it is for this reason that we don't consider an ontology_t as a standard or an agreement: an ontology is not the result of a standards-based concensus of opinion about a portion of the world, because, in general, the effort and thus the result will devolve to the lowest common denominator, and generally end up worthless – because it is inconsistent, has uneven and wrong levels of granularity, and doesn't capture real semantic variances that are crucial for adoption by members of a community. Users of ontology_t cannot be the developers of ontology_t, for much the same reason as users should not develop their own databases: users intuitively know their own semantics, but typically cannot express the ontological and semantic distinctions important to them, nor therefore model them – even though the real world referents are common to everyone.

This avoidance of subjectivism is notwithstanding the established or preferred methodology for developing a specific ontology_t, i.e., that one must focus on the use cases, anticipated scenarios that instantiate those, and therefore the software modeling requirements and "competency questions" (the queries you want answered, i.e., theorems with instantiations which make them true, or the queries you would like to have answered if it were possible for this new ontology-based system to provide you with such), as Fox and Gruninger (1994) and Uschold and Gruninger (1996) clarify. So ontology_t's methodology is to proceed both bottom-up and top-down, i.e., analyze the data sources with respect to their semantics which will have to be captured by the resulting domain ontology and the questions end users (domain experts) would like to ask if they could (and can't typically ask using database and mainstream computing technology). Concerning the latter, typically end users can't formulate these kinds of questions because their imagination is constrained by their current systems. It takes patient knowledge elicitation and knowledge of comparable kinds of value by the working ontology engineer to eke out this kind of knowledge question.

The *focus* of ontological analysis is not centered on the subjective intentions motivating the constructions of the ontology but on the item to be modelled. The reasons for which one is modeling this or that item (or class thereof) may and do interfere, even dramatically, but it is the item itself that is relevant. It is the thing in the world that the ontology is grounded on.

The problem is the old philosophical problem of the connections between epistemology and ontology_c, as we mentioned earlier. The problem has been made obscure by the attitude widely prevalent in recent mainstream philosophy according to which epistemology prevails over ontology.

Ontologies_t could be more or less detailed; their existence may even – at least in some cases – modify the functioning of the modelled system. However, the main question is: does the *existence* of the item/system under observation depend on the ontology? When the answer is negative – as it is for the overwhelming majority of cases – we have a basis for severing the ontological core from all the rest.

The first criterion is then to look for what exists. A number of relatively easy qualifications should now be taken into consideration. The easiest one is to consider only directly observable items, i.e. actually existent items, but also items that existed in the past. They are no more directly observable but we could observe them were we living at their time (This is the pragmatist criterion firstly devised by Peirce. The case where one can *now* observe the traces left by no more existent items is trivially unproblematic). By adopting the same criterion, one may eventually include also items possibly existing in the future.

More demanding is the question about what is said to exist, i.e., the primary interest of ontology_c. For example, we may say that there are material things, plants and animals, as well as the products of the talents and activities of animals and humans in the world. This first prosaic list already indicates that the world comprises not only things, animate or inanimate, but also activities and processes and the products that derive from them. For human-developed products, for example, functional properties are significant (a claw hammer is meant to pound and remove nails; its head and handle are therefore of length and material composition that is appropriately leveragable for those operations). It is likewise difficult to deny that there are thoughts, sensations and decisions, and in fact the entire spectrum of mental activities. Similarly, one is compelled to admit that there are laws, languages, and factories.

We can set about organizing this list of items by saying that there are material items, psychological items and social items (Poli, 2001a), as displayed in Fig. 1.2 below, which depicts dependence of these categories.[17] In turn, each of them presents a vast array of subtypes (material items include physical, chemical, and

[17] Poli's Ontology: The Categorial Stance (TAO-1) discusses these issues in more detail.

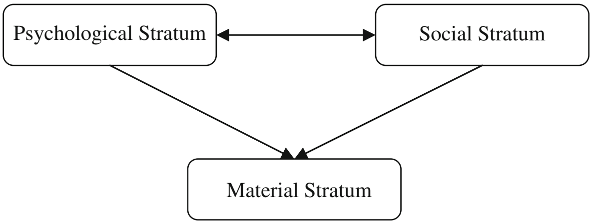

Fig. 1.2 Ontological strata

biological items, psychological items include representation and emotions, social items include laws, languages and many other types of pertinent items).

This section started by asking the natural boundaries of an ontology, how do we determine what an ontology includes or does not include? In trying to provide an answer we found ourselves involved in classical philosophical problems, which is not at all surprising.

1.4.2 Ontology, Science, and Levels of Reality

Returning to our main question, any possibly correct answer concerning what an ontology should include will have to start articulating its proposal with respect to some existing item (or type of). Subsequent steps may follow a variety of different paths. However, for most cases one route seems particularly prominent: that adopted by science. For apparently good reasons, science has been developing in different branches (including physics, economy, biology and cognitive science), the idea being that there are classes of items that "go together", constituting at least a description and possibly an explanation over some portion of reality. In this regard, ontology may follow the same route successfully traversed by science. However different they are, ontology and science are allies. This view intends to convey that between ontology and science there is a mutual exchange, from science to ontology and from ontology to science. That ontology may have something to offer science can be seen from the idea of an ontologically reconstructed and clarified science.

The suggestion is therefore to start from well established scientific partitions. Even if something more will later be required, this initial step will nevertheless help in avoiding two opposed risks. A truly atomistic vision claims that atoms are the only authentically existing items, and that all the other items we recognise are ephemeral. On the other hand, the followers of a boldly holistic vision will claim that the only autonomously existing item is the universe as a whole. Neither of these visions suits ontology. By relying on the multiplicities of sciences one automatically advocates a molar strategy: there are many different items, of different types, requiring different categorial frameworks.

So far so good. The point we arrive at, however, represents both a safe result and one of maximal difficulty. As a matter of fact, so far modern science has relied on an essentially analytic strategy. Different sciences have been developed in order to efficaciously segment the whole of reality into classes of more or less uniformly connected phenomena. The guiding idea has been that phenomena occurring within each class are more causally homogeneous than phenomena pertaining to other classes, so that the task of explaining their behavior should be more easily accomplished. This divide and conquer strategy has proved immensely successful, at least for some regions of reality. Other regions have proved more refractory, for a number of serious reasons. The first is that different regions may require different types of causation, some of which are still unknown, or only partially known (Poli, 2007). A second reason is that for some regions of reality the analytic strategy of breaking items into pieces does not work properly. A third and somewhat connected reason is the lack of a synthetic methodology.

The complexity of reality requires the analytic strategy of segmentation into *categorially* homogeneous regions. This first move is not questioned. However, some regions contain only items that can be further analytically segmented into pieces. These items are entirely governed by their parts (from below, so to speak). Other regions contain items following different patterns: they depend on both their parts and the whole that results from them. Our understanding of these more complex items is still deficient. Recent theories about *granular partitions* (Bittner and Smith, 2001, 2003; Bittner et al, 2007; also see Rogers and Rector, 2000) attempt to remedy this situation.[18] Even so, unfortunately, this is not the end of the story. Something more is further required: sooner or later the products arising from the segmentation into categorially homogeneous regions should be synthesized. For we all live in *one* world. This second synthetic move has proved much more troublesome than the original analytic move.

A properly developed synthetic strategy still awaits us. However, the theories of levels of reality may represent a helpful step toward the elaboration of a fully developed synthetic strategy.[19] Each layer of reality requires (1) specific kinds of items, (2) appropriate categories, and (3) links to its bearing (i.e. based on building-above

[18]Bittner and Smith's (2003) framework tries to uphold the strengths of set theory and mereology for modeling parts and wholes but avoid their respective weaknesses by building on the distinction between *bona fide* (objects which exist independently of human partioning efforts and *fiat* objects (objects which exist only because of human partitioning efforts) (Smith, 2001). As such, their theory of granular partitions begins to impinge on the distinction too between the semantic notions of intension and extension – because on one view, two intensional descriptions ("the morning star", "the evening star") can be seen as human partitions, even though both extensionally refer to the same object, Venus. In their view, "partition is a complex of cells in its projective relation to the world" (Bittner and Smith, 2003, p. 10), and so a triple is established: a granular partition, reality, and the set of "projections" or mappings to and from the items of the partition and reality. Whether this is ontology or ontology intermixed with epistemology remains to be clarified.

[19]Note that we use "level" to refer in general to the levels of reality, restricting the term "layer" to over-forming relationships, and the term "stratum" to building-above relationships. The interested reader is directed to Poli, "Ontology. The Categorial Stance" (TAO-1) for a fuller exposition of this topic.

relations (*Überbauung*), and conditioning (i.e. based on over-forming relations, or *Überformung*)[20] layers as described in Chapter 1, TAO_1.

These are precisely the lacking elements needed to answer the question above asked on the natural boundaries of an ontology: the boundaries of a domain ontology are the top domain categories needed for defining the domain items, plus eventual bearing and conditioning links.

1.4.3 Example: An Ontology of Biology

Using our suggested methodology, we now describe the domain ontology of biology.

Many ontologies_t have been recently developed in the field of biology. Most of them are found at the OBO website.[21] Bio-ontologies offer a nice case for discussing ontological integration. By looking at the ontologies collected by the OBO initiatives one may wonder whether (and how) they could be coordinated and eventually integrated. The problem we would like to address is whether a methodology – here understood as a set of instructions or guidelines – could be devised for developing easy-to-integrate ontologies. Generally speaking, this will comprise minimally three cases: (1) *vertical integration,* i.e., specific ontologies integrable within more general ontologies; e.g. anatomy within biology, (2) *horizontal integration,* i.e., integration among ontologies modelling categorially disjoint phenomena, e.g., business and legal ontologies, and (3) *cross-domain ontologies*, where a number of ontologies pertaining to different levels of reality should be both joined and pruned, e.g., medicine, which may require chemical, biological, psychological, economic, legal and religious information, among others. Some of the mentioned ontologies may further have to be pruned in order to include only information relevant to human pathologies.

The domain top level of our proposed biology ontology will be based on the following three concepts (Hohendorf et al., 2008):

- *Biological entity (BE).*
- *Living entity (LE)*
- *Organism (OR).*

Any biological item is a *biological entity*. The concept of BE refers to anything organic: DNA, mRNA, the nucleus of a cell, the membrane of a cell, its organelles, urine are bioentities. The main function of the concept of BE is to delimit the field. This will prove especially relevant when different ontologies_t are merged together. If all the merged ontologies_t define their most general domain concepts, their management will prove much easier (and safer).

---

[20] Over-forming relations (*Überformung*) and building-above relations (*Überbauung*) are from Hartmann (1952).

[21] Open Biomedical Ontologies (OBO) Foundry. http://obofoundry.org.

Living entities are a specific class of BEs. Two features distinguish LEs from BEs (1) all LEs are systems, while not all BEs are; and (2) LEs are *metabolic* systems. This means that LEs are entities that can survive if appropriate nutrients are provided. Cells, tissues and organs are cases in point.

Lastly come organisms. These are *autopoietic* LEs, i.e. LEs able to produce BEs of which they are composed. Note that a different criterion could be also used: ORs are autonomous LEs (i.e., entities able to survive without a wider encompassing entity). However, we'll use autonomy shortly and will then include it anyway. We have two main types of ORs: unicellular and multicellular ones. The cell that is a unicellular organism is a BE, a LE and an OR. The multicellular organism as a whole is a BE, a LE and an OR. The difference lies in the cell composing the multicellular organism: differently from the single cell that is a unicellular organism, the cells of multicellular organisms are living entities but *are not organisms*. According to the proposed definitions, a liver cell, then, is a living entity but not an organism.

Two more issues remain to be addressed. The first issue concerns the special status of the cell, which sometimes is taken as a full-fledged organism and sometimes not. To mark the difference it is sufficient to consider that all the other living entities apart from cells can *never* be taken as organisms. Cells, therefore, have a special status, that should be marked in some way. We will mark this fact by using the same relation for both the organism-organism case and the cell (as LE) and the organism case. Table 1.1 below shows the relevant cases for binary relations.

Finally, one may notice that if one subtracts LEs from the field of BEs what remains can be taken as coincident with the field of organic chemistry. This is entirely correct and shows the link between a bio-ontology and the ontology/ies characterizing its underlining levels of reality. However, there are differences that shouldn't be forgotten. The most obvious one is that the subtraction of living entities modifies the situation in such a way that a number of question become unanswerable. Even if urine is a purely chemical substance (well, a mixture of), how could one explains its presence without taking organisms into consideration? Organic chemistry and biology present areas of overlapping, and that provides the link between them. However, the questions that are asked from within organic chemistry and the questions that are asked from within biology are different, and this shows that they are different. The reason because they are different lies in the entities grounding their specific levels of reality: molecule for chemistry, (cells and) organisms for biology.

All of this is obviously a first step only. Many other kinds of biological information need to find their proper place. Here is where the second case above mentioned enters the scene. Needless to say, the minimal structure based on the three categories of biological entity, living entity and organism should be widely enlarged if

Table 1.1 Relations between organisms and cells

Relations between unicellular organism A and unicellular organism B.
Relations between unicellular organism A and multicellular organism B (or vice versa).
Relations between cell A (as a LE and not as a OR) and multicellular organism B (or vice versa).

a reasonably articulated model of biology is going to be realized. Once the basic ontological structure of a domain has been established – that is to say, once the levels of reality of the domain have been fixed – the subsequent step is to devising their dimensions of analysis. Here is where faceted analysis can best play its role. Maintaining fixed our reference domain of biology, two series of facets follow. The first series is centered on the governing concept of organism as an individual whole and lists the "viewpoints" from which organisms so taken can be seen. One can then consider at least the following three cases:

- *Classification*
- *Structure*
- *Function*

Classification will model the unquestionably well known biological taxonomies, starting from the distinction between prokaryotes and eucaryotes and then substructuring the former into archea and probacteria and the latter into protista, fungi, plantae, and animalia.

Structure applies part-whole analysis to cells and organisms, for both parts, organ and tissues. Traditionally this type of analysis is called cytology for the cell and anatomy for multicellular organisms. A properly organized biological structure should comprise non only default cases but also variations and anomalies.

Function, finally, corresponds to what is traditionally called physiology and model the working of the part descriptively listed by the previous facet of the organism' structure. In the case of the facet of functions, (serious) variations are usually called pathologies (to be further distinct between intra-systemic and inter-systemic pathologies).

The second series of facets list all the other viewpoints, those not focused on the organism as a whole. These may comprise for instance *genetics* (focus on the genes), *ethology* (focus on some population of organisms), *ecology* (focus on an entire ecosystem). But, again, this is not the entire story. A substantial number of other facets can and should be developed, concerning for instance the *growth* and *development* of organisms, or their *reproduction*, or their *alimentation*. For each of these facets, appropriate ontologies can be developed.

It is time to sum up what has been described in the sections above. We will now try to extract a general scheme from the various topics so far seen.

We propose to distinguish four different types of domain ontologies:

1. *Domain ontologies* in the proper sense (e.g. Biological ontology)
2. *Sub-domain or facet ontologies* (e.g. Gene ontology)
3. *Cross-domain ontologies* (e.g. Medical ontology)
4. *Micro-domain ontologies* (e.g. an ontology of edible substances)

In order to maximize the likelihood of ontology mapping, merging, i.e., integration, we advance two different claims: (1) domain ontologies (of any of the above-distinguished types) should use a top level ontology as their best change of being grounded in a robust framework; (2) furthermore domain ontologies should contain their own top level (domain-top-level); we will further claim that each of the four domain ontologies above distinguished needs a different domain-top-level (this partially explains why distinguishing domain ontologies into specific types is relevant).

In order to distinguish the different cases, the first needed step is to find a criterion for distinguishing domain ontologies in the proper sense from the remaining cases. As a preliminary working definition we propose to define a proper domain ontology as a (1) categorially closed, (2) maximal partition of reality. Therefore, not every partition whatsoever is a domain (in the proper sense).

Our second partition is between domain ontologies (in the proper sense) and their facets. Consider the above described case of a biological ontology. Our claim is that bio-ontology as a whole is a domain ontology while taxonomy, anatomy, physiology, genetics, ecology etc. are some of its facets (and therefore should be properly classified as facet ontologies and not as domain ontologies).

We have claimed that a domain ontology is a categorially autonomous level of reality. “Categorially autonomous” means that even if its phenomena/entities may be existentially dependent from lower and/or side level entities, they nevertheless require a categorial framework different from the one used for understanding the entities of the existentially supporting levels (i.e., they are categorially autonomous).

For instance, biological entities (organisms) require chemical entities (molecules) which in their turn require physical entities (atoms) as their “matter”. However, the frameworks needed for understanding biology, chemistry and physics are different (otherwise there will be only one single science). A somewhat more complex case is provided by side-level domains: the connection linking economy and law is different from the hierarchical one linking say chemistry and biology. The former are domains based on specific phenomena and each requires its specific categorial framework. However, the one supports existentially the other; none of them can exist without the other. Furthermore, both require underlining existential support for both agents and their biological environment.

The above makes clear that one should distinguish existential dependence from categorial autonomy. The latter literally means that the field under analysis presents phenomena requiring new categories.

Any ontology (domain or not) presents a number of different phenomena and usually requires a vast array of categories. Here we claim that all domain ontologies present one (or very few) basic (types of) entities. Sometimes they are difficult to find. Occasionally, science itself has needed quite a while before discovering the constitutive (dominant) entity/entities of the field. The list of domain-constituting entities is an enormously powerful tool for ontology development.

Finally, the top level of a domain ontology should comprise the following information (in parentheses is the information for bio-ontology):

1. The most general domain categories (BE, LE, OR)
2. Link from the domain’s levels of reality (Organism) to the Theory of Levels module of the general top level ontology

3. The namespaces for the domain's "facets" and their position in the overall structure of the domain ontology

The top level of a sub-domain ontology will include only (1) and (3), but may lack (2). (3) may be needed in case the sub-domain ontology is further segmented into sub-sub-domain ontologies.

The last case is the case of cross-domain ontologies, as e.g. medicine. The minimal indication we can here provide is that the top level of a cross-domain ontology should include the top levels of all its relevant domains, plus appropriate pruning of the domain's internal organization. The most obvious case of pruning is limiting biological information to those fragments of biology that are relevant for *homo sapiens*.

1.5 Looking Toward the Future

In this chapter we have demonstrated the interaction between ontology_c and ontology_t, elaborating the definitions of each and describing issues pertinent to both. In addition, we have illustrated this collaboration with an extended example, developing the foundations for an ontology of biology.

Both ontology_c and ontology_t are dependent on each other: (1) ontology_t depends on ontology_c for clarification of the issues involved in describing ontologies about the real world, and (2) ontology_c depends on ontology_t for representing the ontologies about the real world that can be developed into engineering products and be used by software, enabling machines to more closely interact at the human level of conceptualization.

We see increasing interaction between these two aspects of ontology, ontology_c and ontology_t, in the future. Once the more basic issues are elaborated on and scientific consensus established (as we have broached in this chapter), issues such as ontology modularity, mapping, context-determination and representation, vagueness and uncertainty, and ontology lifecycle development will become predominant.

What will become more important are issues ontology_c and ontology_t can work on together and evolve solutions for. Some of these are among the following.

1.5.1 Better Ordering Relations for Ontologies

One issue is the nature of the order relations for ontologies. These can be characterized in many ways: from the perspective of mathematics and computer science set-theoretically, i.e., as partially ordered sets, lattices (semi-lattices; Davey and Priestley, 1991), and including structures used by formal concept analysis (Ganter and Willey, 1996), etc.; or category-theoretically (MacLane, 1971; Lambek and Scott, 1986; Pierce, 1991; Asperti and Longo, 1991; Crole, 1994; Barwise and Seligman, 1997), etc. In ontology_c, these order relations may include those above but extended with notions of mereology and mereotopology (Varzi, 1998), granular partitions (Bitner and Smith, 2003), strata and levels of reality (Poli, 1998, 2001a).

A related issue concerning order is a prospective reconciliation of the notion of subsumption (set-theoretic order) heavily used by informal taxonomies, description logics, thesauri, object-oriented conceptual models, and the logical theories of ontology engineering. Typically this subsumption relation is called *subclass_of* or *isa*, i.e., the special transitive relation of a child class (non-terminal category) to its parent class, with the implicit assumption being made that there is some necessary property that distinquishes a child class from its parent or siblings classes.[22] The subclass/isa relation forms a taxonomic backbone for ontologies_t. Ontology_c has influenced methodologies and tools to assist in the improved development of the taxonomies at the core of ontologies, e.g., the use of meta-properties in OntoClean (Guarino and Welty, 2002).

1.5.2 Elaboration of the Distinctions Among Ontology Levels

We envision increased collaboration in the future between ontology_c and ontology_t on some of the issues we began to elaborate in Section 1.4 focusing on ontology architecture, e.g., the distinctions between upper ontologies and domain ontologies, including levels of reality, and distinct theories of causality (Poli, 2007). These issues also impact the part/whole distinctions made by meronymy, mereology, and topology (Simons, 1988); Varzi and Pianesi, 1996a, b; Varzi, 1998). Discussion of granularity brings up issues related to zooming in and out and reasoning at different levels of reality, approximation, and vagueness (see Williamson, 1998; Keefe and Smith, 1999; Obrst and Mani, 2000; Varzi, 2000; Bittner and Smith, 2001).

Upper ontology issues of ontology_c have already come to the forefront in developing ontology_t products, including time/space distinctions and perspectives: 3D vs. 4D, endurantism vs. perdurantism (see Chapter 14 by Herre, this volume), SNAP/SPAN notions of Grenon (2003), Grenon and Smith (2003). In many cases, ontology engineers working in ontology_t are fiercely supportive of one perspective or another; others would like to maintain the multiple perspectives, picking and choosing which upper theories to use for different domain ontologies or different domain applications. To do so, bridge axiomatization are necessary, e.g., bridging axioms relating 3D to 4D would provide the best of both worlds. However, to date no one has actually created these bridge axioms (Hamlyn, 1984; Sider 2001; Loux, 2002; Obrst et al., 2006).

---

[22]There is also the *instance_of* relation that is the relation between a lowest-level class (non-terminal) or classes (in the case of multiple parents) and the instance (terminal, an individual or particular) which instantiates the properties of that class or classes. In general, classes are universals and instances are particulars.

1.5.3 Ontology Modularity, Mapping, and Formalization of Context

Ontologies_t are often called *logical domain theories* to emphasize that they are logical theories about specific domains of the real world. In the case of upper ontologies, these are *logical foundational (upper) theories*, to signify that they are about notions that traditionally originate from philosophical ontology, ontology_c, but are no less logical theories, i.e., syntactic theories (with licensed semantic models) expressed in a logic and similar to scientific theories except also often extended to common-sense reality (where no scientific theory yet exists or perhaps ever will), and thus a common-sense theory.

Viewed as a collection of interlinked logical theories, ontology_t is concerned with establishing the nature of the relations among these interlinked logical theories, i.e., the nature of the links. Mathematically and computationally, these links are important because they characterize notions of modularity among and within ontologies. If micro-theories can be established (and represented as engineering products) and linked, then these micro-theories can be seen as constituting a theory. Together, many theories can constitute larger theories about reality and on which automated reasoners and other applications in information technology can operate.

It's important that these links among micro-theories and theories are defined logically, so that the "correct entailments" flow across those links. Automated reasoners that reason on vastly many theories require logical notions of modularity, to ensure sound and consistent reasoning. But the definition of what constitutes an ontological "module", i.e., a micro-theory or theory, is not yet agreed on. There are many candidates: one notion of modularity is that of "little theories" in mathematics (Farmer et al., 1992), e.g., how the theory of monoids are related to the theories of groups. There is the microtheory of Cyc (Blair et al., 1992). There is the modularity of category theory (Kent, 2004), which focuses on categories and systems of categories and morphisms among them. There is the so-called "lattice of theories" approach in which distinct ontologies are related by logical relations (or their interpretations) (Sowa, 2000).

But are ontologies_t logical theories about the world or are they vocabularies/models about a community? Some in the information technology community view ontologies as collaborative agreements, more like common conceptual models or standards. This is not a general view, but it exists. Hence, communities of interest form to share information and do so by developing common vocabularies. Typically this is a bottom-up paradigm, wherein communities form to share at least a subset of their individual information, and require common vocabularies and models of those vocabularies. Frequently this paradigm views the process of eliciting these vocabularies and models as a standards activity, wherein concensus is established among a potentially large community.

Directly related to modularity is the the notion of mappings among ontologies, in which disparate ontologies are related. Mapping includes issues such as enforcing local consistency but tolerating global inconsistency (Wiederhold, 1994; Mitra et al., 2000; Noy and Musen, 2000). In general, the mapping between two domain ontologies constitutes a third *integrative domain ontology* which must be able to express everything that is in the two source ontologies.

Finally, pertinent to modularity within and among ontologies is the notion of formalization of context: what does it mean to index assertions of ontologies, to state that certain propositions are true in a given context? Is this the same notion as that of possible worlds with accessibility relations among worlds, with some worlds being "farther away" than other worlds because of their greater respective inconsistencies? Do we need modal logic to express this? Contexts are sometimes called views or perspectives; is this related to the thesaural notion of facets – i.e., perspectives or distinguishable trees connected by cross-references? Context will be increasingly relevant to automated reasoning over ontologies and their knowledge bases. For semantic, computational, and logical approaches to context (see Lewis, 1980; McCarthy, 1987, 1990, 1993; Guha, 1991; McCarthy and Buvač, 1997; Blair et al., 1992; Akman and Surov, 1996, 1997; Giunchiglia and Ghidini, 1998; Menzel, 1999; Bouquet et al., 2003; Obrst et al., 1999a, b; Smith and Obrst, 1999; Obrst and Nichols, 2005; Fikes and Welty, 2006).

1.5.4 Representation vs. Reasoning

This chapter has been primarily focused on the interplay between ontology_c and ontology_t with respect to *representation*, i.e., the correct explication of ontology from the perspective of philosophy and a coherent rendering of that into an engineering model that information technology can utilize. *Representation* here is commonly considered the content and the logical language that the content is expressed in. However, reasoning over that representation is especially important for ontology_t, just as it is for ontology_c. For ontology_c, the reasoning is performed by human beings; for ontology_t, the reasoning is performed by machines. We observe the slogan "No reasoning without representation", which means that automated reasoning, much like human reasoning using formal logical argumentation, is ultimately constrained by the representation that is reasoned over, i.e., the content and the logical language the content is expressed in. One cannot perform full predicate calculus reasoning over content expressed in the propositional calculus, for example.

So a final consideration for future interplay between ontology_c and ontology_t is that focused on the nature of automated reasoning, i.e., the preferred logics for the reasoning, the types of automated reasoning one can perform on those logics (deduction, induction, abduction), the semantics-preserving transformations from one logic or knowledge representation language to another, including perhaps that to special substructural logics (Restall, 2000) for specific kinds of reasoning, and finally, issues of particular relevance to ontology_t concerning *knowledge compilation* (Kautz and Selman, 1994; Cadoli and Donini, 1997; Darwiche and Marquis, 2002), i.e., how best to make an efficient runtime representation for the automated reasoning – which necessarily addresses issues in computational and descriptive complexity (Graedal et al., 2007) from computer science, while at the same time addressing issues in approximation and vagueness on which ontology_c can offer significant insight.

1.5.5 Final Words

Ontology_t needs to be informed by ontology_c and its methods. Ontology_c will increasingly benefit from the sound and consistent software engineering products arising from ontology_t. Computer science, formal philosophy, and formal semantics have come together to birth the beasts called "ontology science" and "ontology engineering". These are strange beasts of Earth but are classifiable and describable under the heavens. Tomorrow it may be that our machines are thereby enabled to in turn be born as stranger beasts, which may yet interact with us human beings as cousins on the well-founded and computationally represented ontological firmament of Earth.

Acknowledgments and Disclaimers

One author's affiliation with The MITRE Corporation is provided for identification purposes only, and is not intended to convey or imply MITRE's concurrence with, or support for, the positions, opinions or viewpoints expressed by this author. We note that the views expressed in this paper are those of the authors alone and do not reflect the official policy or position of any other organization or individual.

References

Akman, V. and S. Mehmet. 1997. The use of situation theory in context modeling. *Computational Intelligence* 13(3):427–438, August, 1997.

Asperti, A. and G. Longo. 1991. *Categories, types and structures*. Cambridge, MA: MIT Press.

Barwise, J., and J. Seligman. 1997. Information flow: The logic of distributed systems. Cambridge, UK: Cambridge University Press.

Basic Formal Ontology (BFO). http://www.ifomis.uni-saarland.de/bfo

Bechhofer, S., F. van Harmelen, J. Hendler, I. Horrocks, D.L. McGuinness, P.F. Patel-Schneider, and L.A. Stein. 2004. OWL Web Ontology Language Reference. W3C Recommendation 10 Feb 2004. http://www.w3.org/TR/owl-ref/

Bittner, T., B. Smith, and M. Donnelly. 2007. The logic of systems of granular partitions. Manuscript. http://ontology.buffalo.edu/smith/articles/BittnerSmithDonnelly.pdf

Bittner, T., and B. Smith. 2003. A theory of granular partitions. In *Foundations of geographic information science*, eds. M. Duckham, M.F. Goodchild, and M.F. Worboys, 117–151, London: Taylor and Francis Books.

Bittner, T., and B. Smith. 2001. Granular partitions and vagueness. In *Formal ontology and information systems*, eds. C. Welty, and B. Smith, 309–321, New York, NY: ACM Press.

Blair, P., R.V. Guha, and W. Pratt. 1992. Microtheories: An ontological engineer's guide. Technical Report Cyc-050-92, 5 Mar 1992, Cycorps, Austin, TX. http://www.cyc.com/tech-reports/cyc-050-92/cyc-050-92.html

Bouquet, P., F. Giunchiglia, F. Van Harmelen, L. Serafini, and H. Stuckenschmidt. 2003. *C-OWL: Contextualizing ontologies*. In 2nd International Semantic Web Conference (ISWC 2003), eds. D. Fensel, K.P. Sycara, and J. Mylopoulos, 164–179, Sanibel Island, FL, 20–23 Oct 2003.

Cadoli, M., and F.M. Donini. 1997. A survey on knowledge compilation. AI communications. *The European Journal for Artificial Intelligence* 10:137–150.

Crole, R.L. 1994. *Categories for types*. Cambridge: Cambridge University Press.

Daconta, M., L. Obrst, K. Smith. 2003. *The semantic web: The future of XML, web services, and knowledge management.* New York, NY: Wiley, June 2003.

Darwiche, A., and P. Marquis. 2001. A knowledge compilation map. http://www.cs.ucla.edu/~darwiche/d116.pdf. An earlier version appeared as "A Perspective on Knowledge Compilation." In Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI'01), Seattle, WA, 175–182.

Davey, B.A., and H.A. Priestley. 1991. *Introduction to lattices and order*. Cambridge, UK: Cambridge University Press.

Descriptive Ontology for Linguistic and Cognitive Engineering (DOLCE). http://www.loa-cnr.it/DOLCE.html

Farmer, W.M., J.D. Guttman, and F.J. Thayer. 1992. Little theories. In *Automated deduction – CADE-11*, LNCS, vol. 607, ed. D. Kapur, 567–581. http://imps.mcmaster.ca/doc/major-imps-papers.html

Fikes, R., and C. Welty. 2006. Interoperable knowledge representation for intelligence support (IKRIS). Advanced Research and Development Activity (ARDA)/Disruptive Technology Office (DTO). Final briefing, Nov 2006.

Fox, M., and M. Gruninger. 1994. *Ontologies for enterprise integration*. In Cooperative Proceedings of the 2nd Conference on Cooperative Information Systems, Toronto, ON.

Ganter, B., and R. Wille. 1996. *Formal concept analysis: Mathematical foundations*. Berlin, Heidelberg, New York: Springer.

General Formal Ontology (GFO). http://www.onto-med.de/en/theories/gfo/index.html

Genesereth, M.R., and Nilsson, N.J. 1987. *Logical foundations of artificial intelligence*. San Mateo, CA: Morgan Kaufmann Publishers.

Graedel, E., P.G. Kolaitis, L. Libkin, M. Marx, J. Spencer, M.Y. Vardi, Y. Venema, and S. Weinstein. 2007. Finite model theory and its applications. Heidelberg: Springer.

Grenon, P. 2003. Spatio-temporality in basic formal ontology: SNAP and SPAN, upper level ontology, and framework of formalization (part I). Technical Report Series 05/2003, IFOMIS.

Grenon, P., and B. Smith. 2003. Snap and span: Towards dynamic geospatial ontology.*Spatial Cognition and Computation* 4(1), forthcoming.

Gruber, T. R. (1991). The role of common ontology in achieving sharable, Reusable knowledge bases. In Principles of Knowledge Representation and Reasoning: Proceedings of the Second International Conference, eds. J.A. Allen, R. Fikes, and E. Sandewall, 601–602, Cambridge, MA: Morgan Kaufmann.

Gruber, T. 1993. A Translation approach to portable ontology specifications. *Knowledge Acquisition* 5:199–220.

Guarino N. 1994. The Ontological Level. Invited Paper Presented at IV Wittgenstein Symposium, Kirchberg, Austria, 1993. In *Philosophy and the cognitive sciences*, eds. R. Casati, B. Smith, and G. White, Vienna: Hölder-Pichler-Tempsky.

Guarino, N, ed. 1998. *Formal ontology and information systems introduction to formal ontology in information systems*. Proceedings of the First International Conference (FOIS'98), 6–8 June 199, Trento, Italy, 3–18. Amsterdam: IOS Press.

Guarino, N., and Giaretta, P. 1995. Ontologies and knowledge bases: Towards a terminological clarification. In *Towards very large knowledge bases: Knowledge building and knowledge sharing*, ed. N. Mars, 25–32, Amsterdam: IOS Press.

Guarino, N., and R. Poli, eds. 1995. Formal ontology in information technology. Special issue of the *International Journal of Human-Computer Studies* 43(5/6). http://www.ladseb.pd.cnr.it/infor/Ontology/IJHCS/IJHCS.html

Guarino, N., and C. Welty. 2002. Evaluating ontological decisions with OntoClean. *Communications of the ACM* 45(2):61–65. New York, NY: ACM Press. http://portal.acm.org/citation.cfm?doid=503124.503150

Guha R.V. 1991. Contexts: A formalization and some applications. PhD Thesis, Stanford University. Also technical report STAN-CS-91-1399-Thesis, and MCC Technical Report Number ACT-CYC-423-91.

Guha, R., and D. Lenat. 1990. Cyc: A mid-term report. Microelectronics Technology and Computer Corporation (MCC), Austin, TX. Technical Report ACT-CYC-134-90.

Hamlyn, D.W. 1984. *Metaphysics*. Cambridge: Cambridge University Press.

Hartmann, N. 1975. *The new ways of ontology.* Westport, Conn.: Greenwood Press.

Herre, H., B. Heller, P. Burek, R. Hoehndorf, F. Loebe, and H. Michalek. 2006. General formal ontology (GFO), part I: Basic principles, Version 1.0. Deliverable No. 8 – July 2006. http://www.onto-med.de/en/publications/scientific-reports/om-report-no8.pdf

Hohendorf, R., F. Loebe, R. Poli, H. Herre, and J. Kelso. 2008. GFO-Bio: A biological core ontology. *Applied Ontology* 3(4):219–227.

Husserl, E. 2001. Logical investigations, vols. 1 and 2, Trans. J.N. Findlay with a new Preface by Michael Dummett, edited with a new introduction by Dermot Moran, Routledge, vols. 1 and 2.

Kautz, H., B. Selman. 1994. An empirical evaluation of knowledge compilation. Proceedings of AAAI-94, Seattle, WA, July 1994.

Kent, R. 2004. The IFF foundation for ontological knowledge organization. In *Knowledge Organization and Classification in International Information Retrieval*, eds. N.J. Williamson, and C. Beghtol, volume 37 of Cataloging and Classification Quarterly, 187–203. New York, NY: Haworth Press.

Interoperable Knowledge Representation for Intelligence Support (IKRIS). http://nrrc.mitre.org/NRRC/ikris.htm

ISO Common Logic. Common logic standard. http://cl.tamu.edu/

Keefe, R., and P. Smith, eds. 1999. *Vagueness: A reader*. Cambridge, MA: MIT Press.

Lewis, D. 1980. Index, context, and content. In *Philosophy and grammar*, eds. S. Kanger, and S. Ohman. Dordrecht: Reidel Publishing.

Lambek, J., and P. Scott. 1986. *Introduction to higher order categorical logic*, volume 7 of Cambridge Studies in Advanced Mathematics. Cambridge: Cambridge University Press.

Loux, M.J. 2002. Metaphysics: A contemporary introduction, 2nd edn. London and New York, NY: Routledge.

Mac Lane, S. 1971. *Categories for the working mathematician*, volume 5 of Graduate Texts in Mathematics. Heidelberg: Springer.

Masolo, C., S. Borgo, A. Gangemi, N. Guarino, and A. Oltramari. 2003. Wonderweb deliverable D18: Ontology library (Final). Technical report, Laboratory for Applied Ontology – ISTC-CNR, Trento.

Margolis, E., and S. Laurence. 1999. *Concepts: Core readings*. Cambridge, MA and London: MIT Press.

McCarthy, J. 1987. Generality in artificial intelligence, *Communications of the ACM* 30(12):1030–1035.

McCarthy, J. 1990. Formalizing common sense: Papers by John McCarthy. Norwood, NJ: Ablex Publishing Corporation.

McCarthy, J. 1993. Notes on formalizing context. In Proceedings of the 13 h International Joint Conference on Artificial Intelligence, Chambery, France.

McCarthy, J., S. Buvač. 1997. Formalizing context (expanded notes). In *Computing natural langauge*, eds. A. Aliseda, R. van Glabbeek, and D. Westerståhl. Stanford, CA: Stanford University. http://www-formal.stanford.edu

Menzel, C. 1999. The objective conception of context and its logic. *Minds and Machines* 9(1):29–56, Feb 1999.

Mitra, P., G. Wiederhold, and M. Kersten. 2000. A graph-oriented model for articulation of ontology interdependencies. Accepted for Extending DataBase Technologies, EDBT 2000, Konstanz, Germany, March 2000. http://www-db.stanford.edu/SKC/publications.html

National Center for Ontological Research (NCOR). http://ncor.buffalo.edu/

Neches, R., R. Fikes, T. Finin, T. Gruber, R. Patil, T. Senator, and W.R. Swartout. 1991. Enabling technology for knowledge sharing. *AI Magazine* 12(3), Fall 1991. http://www.isi.edu/isd/KRSharing/vision/AIMag.html

Noy, N.F., and M.A. Musen. 2000. PROMPT: Algorithm and tool for automated ontology merging and alignment. 17th National Conference on Artificial Intelligence (AAAI-2000), Austin, TX. Technical Report SMI-2000-0831, Stanford Medical Informatics, Stanford University. http://smi-web.stanford.edu/pubs/SMI_Abstracts/SMI-2000-0831.html

Open Biomedical Ontologies (OBO). Foundry. http://obofoundry.org

Object-Centered High-level Reference Ontology (OCHRE). http://www.loa-cnr.it/DOLCE.html

Obrst, L. 2006. What is an ontology? A briefing on the range of semantic models. ontolog forum, 12 and 19 Jan 2006. http://ontolog.cim3.net/cgi-bin/wiki.pl?ConferenceCall_2006_01_12

Obrst, L., T. Hughes, and S. Ray. 2006. Prospects and possibilities for ontology evaluation: The view from NCOR. Workshop on Evaluation of Ontologies for the Web (EON2006), Edinburgh, UK, 22 May 2006.

Obrst, L., and I. Mani, eds. 2000. Proceedings of the Workshop on Semantic Approximation, Granularity, and Vagueness, A Workshop of the Seventh International Conference on Principles of Knowledge Representation and Reasoning KR'2000, Breckenridge, CO, 11 Apr 2000.

Obrst, L, and D. Nichols. 2005. Context and ontologies: Contextual indexing of ontological expressions. AAAI 2005 Workshop on Context and Ontologies, Poster, AAAI 2005, 9–13 July, Pittsburgh, PA. http://www.mitre.org/work/tech_papers/tech_papers_05/05_0903/index.html

Obrst, L., G. Whittaker, A. Meng. 1999a. Semantic context for interoperable distributed object systems. Poster, Modeling and Using Context: Second International and Interdisciplinary Conference (Context'99), Trento, Italy, Sep 1999.

Obrst, L., G. Whittaker, A. Meng. 1999b. Semantic context for object exchange. Workshop on Reasoning in Context for AI Applications, Patrick Brézillon, Roy Turner, Jean-Charles Pomerol, Elise Turner, co-chairs. AAAI-99, Orlando, FL, July, 1999. Technical Report WS-99-14. Menlo Park, CA: AAAI Press.

Ontolog Forum. http://ontolog.cim3.net/cgi-bin/wiki.pl?

Ontology in Information Systems (FOIS-2001), eds. C. Welty, and B. Smith, Ogunquit, Maine, 17–19 Oct 2001.

Ontology Summit. 2007. Ontology, taxonomy, folksonomy: Understanding the distinctions. http://ontolog.cim3.net/cgi-bin/wiki.pl?OntologySummit2007

Petrazycki, L. 1955. *Law and morality*. Partial trans. by H.W. Babb, with an introduction by N.S. Timasheff. Cambridge, MA: Harvard University Press.

Pierce, B. 1991. *Basic category theory for computer scientists*. Cambridge, MA: MIT Press.

Poli, R. Ontology: The categorial stance. See TAO, vol 1.

Poli, R. 1998. Levels. *Axiomathes* 9(1–2):197–211.

Poli, R. 2001a. The basic problem of the theory of levels of reality. *Axiomathes* 12(3–4):261–283.

Poli, R. 2001b. Alwis. Ontology for Knowledge Engineers. PhD Thesis, Utrecht.

Poli, R. 2002. Ontological methodology. *International Journal of Human-Computer Studies* 56:639–664.

Poli, R. 2003. Descriptive, formal and formalized ontologies. In *Husserl's logical investigations reconsidered*, eds. D. Fisette, 193–210. Dordrecht: Kluwer.

Poli, R. 2007. Three obstructions: forms of causation, chronotopoids, and levels of reality. *Axiomathes* 16:1–18.

Ranganathan, S.R. 1962. *Elements of library classification,* 3rd edn. New York, NY: Asia Publishing House.

Restall, G. 2000. *An introduction to substructural logics*. New York, NY and London: Routledge.
Rogers, J.E., and A.L. Rector. 2000. GALEN's model of parts and wholes: Experience and comparisons. Annual Fall Symposium of American Medical Informatics Association, Los Angeles, CA, 714–718. Philadelphia, PA: Hanley and Belfus Inc.
Semy, S., M. Pulvermacher, and L. Obrst. 2005. Toward the use of an upper ontology for U.S. Government and U.S. Military Domains: An evaluation. MITRE Technical Report, MTR 04B0000063, Nov 2005. http://www.mitre.org/work/tech_papers/tech_papers_05/04_1175/index.html
Sider, T. 2001. *Four-dimensionalism. An ontology of persistence and Time*. Oxford: Clarendon Press.
Simons, P. 1987. *Parts: A study in ontology*. Oxford: Clarendon Press.
Smith, B. 2001a. Fiat objects. *Topoi* 20(2):131–148.
Smith, K., L. Obrst. 1999. Unpacking the semantics of source and usage to perform semantic reconciliation in large-scale information systems. SIGMOD special issue on Semantic Interoperability, eds. A. Sheth, and A. Ouksel, SIGMOD, Mar 1999.
Sowa, J. 2000. *Knowledge representation: Logical, philosophical, and computational foundations*. Pacific Grove, CA: Brooks/Cole Thomson Learning.
Suggested Upper Merged Ontology (SUMO). http://www.ontologyportal.org/
Upper Cyc. http://www.cyc.com/cycdoc/vocab/vocab-toc.html
Uschold, M., M. Gruninger. 1996. Ontologies: Principles, methods, and applications. *The Knowledge Engineering Review* 11(2):93–136.
Varzi, A. 2000. Vagueness, logic, and ontology, to appear in The Dialogue. http://www.columbia.edu/~av72/papers/Dialogue_2000.pdf
Varzi, A.C. 1998. Basic problems of mereotopology. In *Formal ontology in information systems*, eds. N. Guarino, 29–38. Amsterdam: IOS Press.
Varzi, A., and F. Pianesi. 1996a. Events, topology, and temporal relations. *The Monist* 78(1): 89–116.
Varzi, A., and F. Pianesi. 1996b. Refining temporal reference in event structures. *Notre Dame Journal of Formal Logic* 37(1):71–83.
Wiederhold, G. 1994. An algebra for ontology composition. Proceedings of 1994 Monterey Workshop on Formal Methods, Sept 1994, U.S. Naval Postgraduate School, Monterey, CA, 56–61. http://www-db.stanford.edu/pub/gio/paperlist.html
Williamson, T. 1998. *Vagueness*. London and New York, NY: Routledge.