Ontology

Barry Smith

Preprint version of a paper forthcoming in Luciano Floridi (ed.), *The Blackwell Guide to the Philosophy of Information and Computing*, Oxford: Blackwells, forthcoming.

Philosophical Ontology

Ontology as a branch of philosophy is the science of what is, of the kinds and structures of objects, properties, events, processes and relations in every area of reality. 'Ontology' is often used by philosophers as a synonym of 'metaphysics' (a label meaning literally: 'what comes after the *Physics*'), a term used by early students of Aristotle to refer to what Aristotle himself called 'first philosophy'. Sometimes 'ontology' is used in a broader sense, to refer to the study of what *might* exist; 'metaphysics' is then used for the study of which of the various alternative possible ontologies is in fact true of reality. (Ingarden 1964) The term 'ontology' (or *ontologia*) was coined in 1613, independently, by two philosophers, Rudolf Göckel (Goclenius), in his *Lexicon philosophicum* and Jacob Lorhard (Lorhardus), in his *Theatrum philosophicum*. Its first occurrence in English as recorded by the OED appears in Bailey's dictionary of 1721, which defines ontology as 'an Account of being in the Abstract'.

Ontology seeks to provide a definitive and exhaustive classification of entities in all spheres of being. The classification should be definitive in the sense that it can serve as an answer to such questions as: What classes of entities are needed for a complete description and explanation of all the goings-on in the universe? Or: What classes of entities are needed to give an account of what makes true all truths? It should be exhaustive in the sense that all types of entities should be included in the classification, including also the types of relations by which entities are tied together to form larger wholes.

Different schools of philosophy offer different approaches to the provision of such classifications. One large division is that between what we might call substantialists and fluxists, which is to say between those who conceive ontology as a substance- or thing- (or continuant-) based discipline and those who favour an ontology centred on events or processes (or occurrents). Another large division is between what we might call adequatists and reductionists. Adequatists seek a taxonomy of the entities in reality at all levels of aggregation, from the microphysical to the cosmological, and including also the middle world (the *mesocosmos*) of human-scale entities in between. Reductionists see reality in terms of some one privileged level of existents; they seek to establish the 'ultimate furniture of the universe' by decomposing reality into its simplest constituents, or they seek to 'reduce' in some other way the apparent variety of types of entities existing in reality.

It is the work of adequatist philosophical ontologists such as Aristotle, Ingarden (1964), and Chisholm (1996) which will be of primary importance for us here. Their taxonomies are in many ways comparable to the taxonomies produced by sciences such as biology or chemistry, though they are of course radically more general than these. Adequatists transcend the dichotomy between substantialism and fluxism, since they accept categories of both continuants and occurrents. They study the totality of those objects, properties, processes and relations that make up the world on different levels of focus and granularity, and whose different parts and moments are studied by the different scientific disciplines. Ontology, for the adequatist, is then a descriptive enterprise. It is thus distinguished from the special sciences not only in its radical generality but also in its goal or focus: it seeks not predication, but rather taxonomy.

Methods of Ontology

The methods of ontology – henceforth in philosophical contexts always used in the adequatist sense – are the methods of philosophy in general. They include the development of theories of wider or narrower scope and the testing and refinement of such theories by measuring them up, either against difficult counterexamples or against the results of science. These methods were familiar already to Aristotle himself.

In the course of the twentieth century a range of new formal tools became available to ontologists for the development and testing of their theories. Ontologists nowadays have a choice of formal frameworks (deriving from algebra, category theory, mereology, set theory, topology) in terms of which their theories can be formulated. These new formal tools, along with the language of formal logic, allow philosophers to express intuitive principles and definitions in clear and rigorous fashion, and, through the application of the methods of formal semantics, they can allow also for the testing of theories for consistency and completeness.

Ontological Commitment

To create effective representations it is an advantage if one knows something about the things and processes one is trying to represent. (We might call this *the Ontologist's Credo.*) The attempt to satisfy this credo has led philosophers to be maximally opportunistic in the sources they have drawn upon in their ontological explorations of reality and in their ontological theorizing. These have ranged all the way from the preparation of commentaries on ancient texts to reflection on our linguistic usages when talking about entities in domains of different types. Increasingly, however, philosophers have turned to science, embracing the assumption that one (perhaps the only) generally reliable way to find out something about the things and processes within a given domain is to see what scientists say. Some philosophers have thought that the way to do ontology is exclusively through the investigation of scientific theories.

With the work of Quine (1953) there arose in this connection a new conception of the proper method of ontology according to which the ontologist's task is to establish what kinds of entities scientists are committed to in their theorizing. The ontologist studies the world by drawing conclusions from the theories of the natural sciences, which Quine takes to be our best sources of knowledge as to what the world is like. Such theories are extensions of the theories we develop and use informally in everyday life, but they are developed with closer attention to certain special kinds of evidence that confer a higher degree of probability on the claims made. Quine takes ontology seriously. His aim is to use science for ontological purposes, which means: to find the ontology *in* scientific theories. Ontology is then a network of claims, derived from the natural sciences, about what exists coupled with the attempt to establish what types of entities are most basic. Each natural science has, Quine holds, its own preferred repertoire of types of objects to the existence of which it is committed. Each such theory embodies only a partial ontology. This is defined by the vocabulary of the corresponding theory and (most importantly for Quine) by its canonical formalization in the language of first-order logic.

Note that ontology is for Quine himself not the meta-level study of the ontological commitments or presuppositions embodied in the different natural-scientific theories. Ontology is rather these commitments themselves. Quine moves to the meta-level, making a semantic ascent to consider the statements in a theory, only in setting out to establish those expressions which definitively carry its commitments. Quine fixes upon the language of first-order logic as the medium of canonical representation not out of dogmatic devotion to this particular form, but rather because he holds that this is the only really clear form of language. First-order logic is itself just a regimentation of corresponding parts of ordinary language, a regimentation from which, in Quine's eyes, logically problematic features have been excised. It is then, Quine argues, only the bound variables of a theory that carry its definitive commitment to existence. It is sentences like 'There are horses,' 'There are numbers,' 'There are electrons,' that do this job. His so-called 'criterion of ontological commitment' is captured in the slogan: *To be is to be the value of a bound variable.* This should not be understood as signifying some reductivistic conception of existence itself as a merely logico-linguistic matter. Rather it is to be interpreted in practical terms: to determine what the ontological commitments of a scientific theory are, it is necessary to determine the values of the quantified variables used in its canonical formalizations.

Quine's approach is thus most properly conceived not as a reduction of ontology to the study of scientific language, but rather as a continuation of ontology in the traditional sense. When viewed in this light, however, it can be seen to be in need of vital supplementation. For the objects of scientific theories are discipline-specific. This means that the *relations* between objects belonging to different disciplinary domains fall out of bounds for Quinean ontology. Only something like a *philosophical* theory of how different scientific theories (or their objects) relate to each other can fulfil the task of providing an inventory of all the types of entities in reality. Quine himself would resist this latter conclusion. For him the best we can achieve in ontology lies in the quantified statements of particular theories, theories supported by the best evidence we can muster. We have no way to rise above the particular theories we have; no way to harmonize and unify their respective claims.

Internal vs. External Metaphysics

Quine is a realist philosopher. He believes in a world beyond language and beliefs, a world which the theories of natural science give us the power to illuminate. There is, however, another tendency in twentieth-century analytic philosophy, a tendency often associated with Quine but inspired much rather by Kant and promulgated by thinkers such as Carnap and Putnam, according to which ontology is a meta-level discipline which concerns itself not with the world itself but rather only with theories or languages or systems of beliefs. Ontology as a first-level discipline of the world beyond – ontology as what these philosophers call 'external metaphysics' – is impossible. The best we can achieve, they hold, is *internal metaphysics*, which means precisely the study of the ontological commitments of specific theories or systems of beliefs. Strawsonian descriptive metaphysics is one example of such internal metaphysics. Model-theoretic semantics, too, is often implicitly understood in internal-metaphysical terms – the idea being that we cannot understand what a given language or theory is really about, but we can build *models* with more or less nice properties. What we can never do is compare these models to some reality beyond. Ontology in the traditional philosophical sense thus comes to be replaced by the study of how a given language or science conceptualizes a given domain. It becomes a theory of the ontological content of certain representations. Traditional ontologists are seeking principles that are true of reality. The practitioners of internal metaphysics, in contrast, are seeking to elicit principles from subjects or theories. The elicited principles may or may not be true, but this, to the practitioner of internal metaphysics, is of no concern, since the significance of these principles lies elsewhere – for instance in yielding a correct account of the taxonomical system used by speakers of a given language or by the scientists working in a given discipline.

In a development that has hardly been noted by philosophers, a conception of the job of the ontologist close to that of Carnap and Putnam has been advanced in recent years also in certain extra-philosophical disciplines, as linguists, psychologists and anthropologists have sought to elicit the ontological commitments ('ontologies', in the plural) of different cultures and groups. Thus, they have sought to establish the ontology underlying common-sense or folk theories of various sorts by using the standard empirical methods of the cognitive sciences (see for example Keil 1979, Spelke 1990). Researchers in psychology and anthropology have sought to establish what individual human subjects, or entire human cultures, are committed to, ontologically, in their everyday cognition, in much the same way in which Quine-inspired philosophers of science had attempted to elicit the ontological commitments of the natural sciences.

It was still reasonable for Quine to identify the study of ontology – the search for answers to the question: what exists? – with the study of the ontological commitments of natural scientists. This is because it is a reasonable hypothesis that all natural sciences are in large degree consistent with each other. Moreover, the identification of ontology with ontological commitments continues to seem reasonable when one takes into account not only the natural sciences but also certain commonly shared commitments of common sense – for example that *tables* and *chairs* and *people* exist. For the common-sense taxonomies of objects such as these are compatible with those of scientific theory if only we are careful to take into account the different granularities at which each operates. (Forguson 1989, Omnès 1999, Smith and Brogaard 2002)

Crucially, however, the identification of ontology with ontological commitments becomes strikingly less defensible when the ontological commitments of various specialist groups of *non*-scientists are allowed into the mix. How, ontologically, are we to treat the commitments of astrologists, or clairvoyants, or believers in leprechauns? We shall return to this question below.

Ontology and Information Science

In a related development, also hardly noticed by philosophers, the term 'ontology' has gained currency in recent years in the field of computer and information science.

The first big task for the new 'ontology' derives from what we might call the Tower of Babel problem. Different groups of data- and knowledge-base system designers have their own idiosyncratic terms and concepts by means of which they build frameworks for information representation. Different databases may use identical labels but with different meanings; alternatively the same meaning may be expressed via different names. As ever more diverse groups are involved in sharing and translating ever more diverse varieties of information, the problems standing in the way of putting this information together within a single system increase geometrically. Methods must be found to resolve the terminological and conceptual incompatibilities which then inevitably arise.

Initially, such incompatibilities were resolved on a case-by-case basis. Gradually, however, it was recognized that the provision, once and for all, of a common reference ontology – a shared taxonomy of entities – might provide significant advantages over such case-by-case resolution, and the term 'ontology' came to be used by information scientists to describe the construction of a canonical description of this sort. An ontology is in this context a dictionary of terms formulated in a canonical syntax and with commonly accepted definitions designed to yield a lexical or taxonomical framework for knowledge-representation which can be shared by different information systems communities. More ambitiously, an ontology is a formal theory within which not only definitions but also a supporting framework of axioms is included (perhaps the axioms themselves provide implicit definitions of the terms involved).

The methods used in the construction of ontologies thus conceived are derived on the one hand from earlier initiatives in database management systems. But they also include methods similar to those employed in philosophy (as described in Hayes 1985), including the methods used by logicians when developing formal semantic theories.

Upper-Level Ontologies

The potential advantages of ontology for the purposes of information management are obvious. Each group of data analysts would need to perform the task of making its terms and concepts compatible with those of other such groups only once – by calibrating its results in the terms of the single canonical backbone language. If all databases were calibrated in terms of just one common ontology (a single consistent, stable and highly expressive set of category labels), then the prospect would arise of leveraging the thousands of person-years of effort that have been invested in creating separate database resources in such a way as to create, in more or less automatic fashion, a single integrated knowledge base of a scale hitherto unimagined, thus fulfilling an ancient philosophical dream of a Great Encyclopedia comprehending all knowledge within a single system.

The obstacles standing in the way of the construction of a single shared ontology in the sense described are unfortunately prodigious. Consider the task of establishing a common ontology of world history. This would require a neutral and common framework for all descriptions of historical facts, which would require in turn that all legal and political systems, rights, beliefs, powers, and so forth, be comprehended within a single, perspicuous list of categories.

Added to this are the difficulties which arise at the level of adoption. To be widely accepted an ontology must be neutral as between different data communities, and there is, as experience has shown, a formidable trade-off between this constraint of neutrality and the requirement that an ontology be maximally wide-ranging and expressively powerful – that it should contain canonical definitions for the largest possible number of terms. One solution to this trade-off problem is the idea of a top-level ontology, which would confine itself to the specification of such highly general (domain-independent) categories as: time, space, inherence, instantiation, identity, measure, quantity, functional dependence, process, event, attribute, boundary, and so on. (See for example http://suo.ieee.org.) The top-level ontology would then be designed to serve as common neutral backbone, which would be supplemented by the work of ontologists working in more specialized domains on, for example, ontologies of geography, or medicine, or ecology, or law, or, still more specifically, ontologies of built environments (Bittner 2001), or of surgical deeds (Rossi

Mori *et al.* 1997).

Uses of Ontology

The initial project of building one single ontology, even one single top-level ontology, which would be at the same time non-trivial and also readily adopted by a broad population of different information systems communities, has largely been abandoned. The reasons for this can be summarized as follows. The task of ontology-building proved much more difficult than had initially been anticipated (the difficulties being at least in part identical to those with which philosophical ontologists have grappled for some 2000 years). The information systems world itself, on the other hand, is very often subject to the short time horizons of the commercial environment. This means that the requirements placed on information systems change at a rapid rate, so that already for this reason work on the construction of corresponding ontological translation modules has been unable to keep pace.

Yet work in ontology in the information systems world continues to flourish, and the principal reason for this lies in the fact that its focus on classification (on analysis of object types) and on constraints on allowable taxonomies has proved useful in ways not foreseen by its initial progenitors (Guarino and Welty 2000). The attempt to develop terminological standards, which means the provision of explicit specifications of the meanings of the terms used in application domains such as medicine or air traffic control, loses nothing of its urgency even when it is known in advance that the more ambitious goal of a common universal ontology is unlikely to be realized.

Consider the following example, due to Guarino. Financial statements may be prepared under either the US GAAP or the IASC standards (the latter being applied in Europe and many other countries). Under the two standards, cost items are often allocated to different revenue and expenditure categories depending on the tax laws and accounting rules of the countries involved. So far it has not been possible to develop an algorithm for the automatic conversion of income statements and balance sheets between the two systems, since so much depends on highly volatile case law and on the subjective interpretation of accountants. Not even this relatively simple problem has been satisfactorily resolved, though this is *prima facie* precisely the sort of topic where ontology could contribute something of great commercial impact.

If Ontek did not Exist, it would be Necessary to Invent it.

Perhaps the most impressive attempt to develop an ontology – at least in terms of sheer size – is the CYC project (http://www.cyc.com), which grew out of an effort initiated by Doug Lenat in the early 1980s to formalize common-sense knowledge in the form of a massive database of axioms covering all things, from governments to mothers. The resultant ontology has been criticised for its ad hoc (which is to say: unprincipled) nature. It takes the form of a tangled hierarchy, with a topmost node labelled *Thing*, beneath which are a series of cross-cutting total partitions, including: *Represented Thing* vs. *Internal Machine Thing*, *Individual Object* vs. *Collection*, *Intangible* vs. *Tangible Object* vs. *Composite Tangible and Intangible Object.* Examples of Intangible Objects (*Intangible* means: *has no mass*) are sets and numbers. A person, in the CYC ontology, is a Composite Object made up of a Tangible body and an Intangible mind.

More important, for our purposes here, however, is the work of the firm Ontek – short for 'ontological technology' – which since 1981 has been developing database programming and knowledge representation technologies necessary to create decision automation systems – "white collar robots" – for large-scale industrial enterprises in fields such as aerospace and defense. Realizing that the ontology required to build such systems would need to embrace in a principled way the entire gamut of entities encompassed by these businesses in a single, unified framework, Ontek approached this problem by systematically exploiting the resources of ontology in the traditional (adequatist) philosophical sense. A team of philosophers (including David W. Smith and Peter Simons) collaborated with software engineers in constructing the system PACIS (for Platform for the Automated Construction of Intelligent Systems), which is designed to implement a comprehensive theory of entities, ranging from the very concrete (aircraft, their structures, and the processes involved in designing and developing them) to the somewhat abstract (business processes and organizations, their structures, and the strategies involved in creating them) to the exceedingly abstract formal structures which bring all of these diverse components together.

Ontek has thus realized in large degree the project sketched by Hayes in his " Naïve Physics Manifesto", of building a massive theory of (in Hayes' case) common-sense physical reality (in Ontek's case this is extended to include not only airplane wings and factories but also associated planning and accounting procedures), by putting away the toy worlds of classical AI research and concentrating instead on the formalization of the ontological features of the world as this is encountered by adults engaged in the serious business of living. Such large-scale projects are, as Hayes already recognized, essential for long-term progress in artificial intelligence. Where Hayes conceived his project as that of formalizing our 'mental models' – so that his "Naïve Physics Manifesto", like Lenat's CYC, is a contribution not to the discipline of ontology in the traditional sense at all, but rather to that of knowledge representation – Dement and his collaborators have taken the bull of reality by the horns, and sought to develop a true theory of the world from the vantage point of large-scale commercial enterprises.

The Leipzig project in medical ontology (see http://ifomis.de and Degen *et al.* 2001), too, is based on a realist methodology close to that of Ontek. Most prominent information systems ontologists in recent years, however, have abandoned the Ontologist's Credo and have embraced instead a view of ontology as an inwardly directed discipline (so that they have in a sense adopted an epistemologized reading of ontology itself analogous to that of Carnap and Putnam). They have come to hold, with Gruber (1995), that: 'For AI systems what "exists" is that which can be represented.' This means not only that only those entities exist which are represented in the system, but also that such entities can possess only those properties which are represented in the system. It is as if Hamlet, whose hair (we shall suppose) is not mentioned in Shakespeare's play, would be not merely neither bald nor non-bald, but would somehow *have no properties at all* as far as hair is concerned. (Compare Ingarden (1973) on the 'loci of indeterminacy' within the stratum of represented objects of a literary work.) What this means, however, is that the objects represented in the system (for example people in a database) are not real objects – the objects of flesh and blood we find all around us. Rather, they are denatured surrogates, possessing only a finite number of properties (sex, date of birth, social security number, marital status, employment status, and the like), and being otherwise entirely indeterminate with regard to all those properties and dimensions with which the system is not concerned.

Information systems ontologies in the sense of Gruber are, we see, not oriented around the world of objects at all. Rather, they are focused on our concepts or languages or mental models (or, on a less charitable interpretation, objects and concepts are simply confused). It is in this light that we are to interpret passages such as the following:

> an ontology is a description (like a formal specification of a program) of the concepts and relationships that can exist for an agent or a community of agents. This definition is consistent with the usage of ontology as set-of-concept-definitions, but more general.

> And it is certainly a different sense of the word than its use in philosophy. (Gruber, n.d.)

Conceptualizations

The newly fashionable usage of 'ontology' as meaning just 'conceptual model' is by now firmly entrenched in many information systems circles. Gruber is to be given credit for having crystallized the new sense of the term by relating it to the technical definition of 'conceptualization' introduced by Genesereth and Nilsson in their (1987). In his (1993) Gruber defines an ontology as 'the specification of a conceptualization'. Genesereth and Nilsson conceive conceptualisations as extensional entities (they are defined in terms of sets of relations), and they have accordingly been criticized on the grounds that this extensional understanding makes conceptualizations too remote from natural language, where intensional contexts predominate (see Guarino, Introduction to 1998). For present purposes, however, we can ignore these issues, since we shall gain a sufficiently precise understanding of the nature of 'ontology', as Gruber conceives it, if we rely simply on the account he himself gives in passages such as the following:

> A conceptualization is an abstract, simplified view of the world that we wish to represent for some purpose. Every knowledge base, knowledge-based system, or knowledge-level agent is committed to some conceptualization, explicitly or implicitly. (Gruber 1995)

The idea is as follows. As we engage with the world from day to day we participate in rituals and we tell stories. We use information systems, databases, specialized languages, and scientific instruments. We buy insurance, negotiate traffic, invest in bond derivatives, make supplications to the gods of our ancestors. Each of these ways of behaving involves, we can say, a certain conceptualization. What this means is that it involves a system of concepts in terms of which the corresponding universe of discourse is divided up into objects, processes and relations in different sorts of ways. Thus in a religious ritual setting we might use concepts such as *salvation* and *purification*; in a scientific setting we might use concepts such as *virus* and *nitrous oxide*; in a story-telling setting we might use concepts such as: *leprechaun* and *dragon*. Such conceptualizations are often tacit; that is, they are often not thematized in any systematic way. But tools can be developed to specify and to clarify the concepts involved and to establish their logical structure, and thus to render explicit the underlying taxonomy. We get very close to the use of the term 'ontology' in Gruber's sense if we define an ontology as the result of such clarification – as, precisely, the *specification of a conceptualization* in the intuitive sense described in the above.

Ontology thus concerns itself not at all with the question of ontological realism, that is with the question whether its conceptualizations are *true of* some independently existing reality. Rather, it is a strictly pragmatic enterprise. It starts with conceptualizations, and goes from there to the description of corresponding domains of objects (often confusingly referred to as 'concepts'), but the latter are nothing more than nodes in or elements of closed world data models devised with specific practical purposes in mind.

What is most important, now, is that all of the mentioned surrogate created worlds are treated by the ontological engineer as being on an equal footing. In a typical case the universe of discourse will be specified by the client or customer, and for the purposes of the ontological engineer the customer is always right (it is the customer in each case who defines his own specific world of surrogate objects). It is for this reason that the ontological engineer aims not for truth, but rather, merely, for adequacy to whatever is the pertinent application domain as defined by the client. The main focus is on reusability of application domain knowledge in such a way as to accelerate the development of similar software systems in each new application context. The goal is not truth to some independently existing domain of reality, which is after all often hard to achieve, but merely (at best) truth relative to some conceptualisation.

Given this background we can understand why the project of a common ontology which would be accepted by many different information communities in many different domains has thus far failed. Not all conceptualizations are equal. What the customer says is not always true; indeed it is not always sufficiently coherent to be even in the market for being true. Bad conceptualizations abound (rooted in error, myth-making, astrological prophecy, hype, bad linguistics, or antiquated information systems based on dubious foundations). Such conceptualisations may deal *only* with created (pseudo-)domains, and not with any transcendent reality beyond.

Consider, now, against this background the project of developing a top-level ontology, a common ontological backbone. It begins to seem rather like the attempt to find some highest common denominator that would be shared in common by a plurality of true and false theories. Seen in this light, the principal reason for the failure of attempts thus far to construct top-level ontologies lies precisely in the fact that these attempts were made on the basis of a methodology which treated all application domains on an equal footing. It thereby overlooked the degree to which the different conceptualizations which serve as inputs to ontology are likely to be not only of wildly differing quality but also mutually inconsistent.

What can Information Scientists learn from Philosophical Ontologists?

As we have seen, some ontological engineers have recognized that they can improve their models by drawing on the results of the philosophical work in ontology carried out over the last 2000 years. This does not in every case mean that they are ready to abandon their pragmatic perspective. Rather, they see it as useful to employ a wider repertoire of ontological theories and frameworks and, like philosophers themselves, they are willing to be maximally opportunistic in their selection of resources for purposes of ontology-construction. Guarino and his collaborators, for example, use standard philosophical analyses of notions such as identity, set-theoretical subsumption, part-whole subsumption and the like in order to expose inconsistencies in standard upper-level ontologies such as CYC, and they go on from there to derive meta-level constraints which all ontologies must satisfy if they are to avoid inconsistencies of the sorts exposed.

Given what was said above, however, it appears that information ontologists may have sound *pragmatic* reasons to take the philosopher ontologist’s traditional concern for truth more seriously still. For the very abandonment of the focus on mere conceptualisations and on conceptualisation-generated object-surrogates may itself have positive pragmatic consequences.

This applies even in the world of administrative systems, for example in relation to the GAAP/IASC integration problem referred to above. For ontologists are here working in a type of theoretical context where they must move back and forth between distinct conceptualisations, and where they can find the means to link the two together only by looking at their common objects of reference in the real, flesh-and-blood world of human agents and financial transactions.

Where ontology is directed in this fashion, not towards a variety of more or less coherent surrogate models, but rather towards the real world of flesh-and-blood objects in which we all live, then this itself reduces the likelihood of inconsistency and systematic error in the theories which result, and, conversely, it increases the likelihood of our being able to build a single workable system of ontology that will be at the same time non-trivial. On the other hand, however, the ontological project thus conceived will take much longer to complete and it will face considerable internal difficulties along the way. Traditional ontology is a difficult business. At the same time, however, it has the potential to reap considerable rewards – not least in terms of a greater stability and conceptual coherence of the software artefacts constructed on its basis.

To put the point another way: it is precisely because good conceptualizations are transparent to reality that they have a reasonable chance of being integrated together in robust fashion into a single unitary ontological system. The fact that the real world itself plays a significant role in ensuring the unifiability of our separate ontologies thus implies that, if we are to accept a conceptualization-based methodology as one stepping stone towards the construction of adequate ontologies, then we must abandon the attitude of tolerance towards both good and bad conceptualizations. For it is this very tolerance which is fated to undermine the project of ontology itself.

Of course to zero in on good conceptualizations is no easy matter. There is no Geiger-counter-like device which can be used for automatically detecting truth. Rather, we have to rely at any give stage on our best endeavors – which means concentrating above all on the work of natural scientists – and proceed in careful, critical and fallibilistic fashion from there, hoping to move gradually closer to the truth via an incremental process of theory construction, criticism, testing, and amendment. As suggested in Smith and Mark (2001) there may be reasons to look beyond natural science, above all where we are dealing with objects (such as societies, institutions and concrete and abstract artefacts) existing at levels of granularity distinct from those which readily lend themselves to natural-scientific inquiry. Our best candidates for good conceptualizations will however remain those of the natural sciences – so that we are, in a sense, brought back to Quine, for whom the job of the ontologist coincides with the task of establishing the ontological commitments of scientists, and of scientists alone.

What Can Philosophers Learn from Information Systems Ontologists?

Developments in modal, temporal and dynamic logics as also in linear, substructural and paraconsistent logics have demonstrated the degree to which advances in computer science can yield benefits in logic – benefits not only of a strictly technical nature, but also sometimes of wider philosophical significance. Something similar can be true, I suggest, in relation to the developments in ontological engineering referred to above. The example of the successes and failures of information systems ontologists can first of all help to encourage existing tendencies in philosophical ontology (nowadays often grouped under the heading 'analytic metaphysics') towards opening up new domains of investigation, for example the domain of social institutions (Mulligan 1987, Searle 1995), of patterns (Johansson 1998), of artefacts (Dipert 1993, Simons and Dement 1996), of boundaries (Smith 2001), of dependence and instantiation (Mertz 1996, Degen *et al*., 2001), of holes (Casati and Varzi 1994), and parts (Simons 1987). Secondly, it can shed new light on the many existing contributions to ontology, from Aristotle to Goclenius and beyond (Burkhardt and Smith 1991), whose significance was for a long time neglected by philosophers in the shadow of Kant and other enemies of metaphysics. Thirdly, if philosophical ontology can properly be conceived as a kind of generalized chemistry, then information systems can help to fill one important gap in ontology as it has been practiced thus far, which lies in the absence of any analogue of chemical experimentation. For one can, as C. S. Peirce remarked (1933, 4.530), 'make exact experiments upon uniform diagrams'. The new tools of ontological engineering might help us to realize Peirce's vision of a time when operations upon diagrams will 'take the place of the experiments upon real things that one performs in chemical and physical research.'

Finally, the lessons drawn from information systems ontology can support the efforts of those philosophers who have concerned themselves not only with the development of ontological *theories*, but also – in a field sometimes called 'applied ontology' (Koepsell 1999, 2000) – with the *application* of such theories in domains such as law, or commerce, or medicine. The tools of philosophical ontology have been applied to solve practical problems, for example concerning the nature of intellectual property or concerning the classification of the human foetus at different stages of its development. Collaboration with information systems ontologists can support such ventures in a variety of ways, first of all because the results achieved in specific application-domains can provide stimulation for philosophers, but also – and not least importantly – because information systems ontology is itself an enormous new field of practical application that is crying out to be explored by the methods of rigorous philosophy.

Acknowledgements

This chapter is based upon work supported by the National Science Foundation under Grant No. BCS-9975557 ("Ontology and Geographic Categories"). Thanks go to the NSF and to Thomas Bittner, Charles Dement, Andrew Frank, Angelika Franzke, Wolfgang Grassl, Nicola Guarino, Kathleen Hornsby, Ingvar Johansson, Kevin Mulligan, David W. Smith, William Rapaport, Chris Welty and Graham White for helpful comments. They are not responsible for any errors which remain.

Literature

Bittner, Thomas 2001 "The Qualitative Structure of Built Environments," *Fundamenta Informaticae*, 46, 97–126. Uses the theory of fiat boundaries to develop an ontology of urban environments.

Brentano, Franz 1981 *The Theory of Categories*, The Hague/Boston/London: Martinus Nijhoff. Defends a classification of entities, and a new mereological view of substances and their accidents, based on Aristotle.

Burkhardt, Hans and Smith, Barry (eds.) 1991 *Handbook of Metaphysics and Ontology*, 2 vols., Munich/Philadelphia/Vienna: Philosophia. Reference work on philosophical ontology and ontologists.

Casati, Roberto and Varzi, Achille C. 1994 *Holes and Other Superficialities*, Cambridge, Mass.: MIT Press. On the ontology and cognition of holes.

Chisholm, Roderick M. 1996 *A Realistic Theory of Categories: An Essay on Ontology*, Cambridge: Cambridge University Press. Defends a classification of entities based on Aristotle and Brentano.

Degen, W., Heller, B., Herre, H. and Smith, B. 2001 "GOL: Towards an Axiomatized Upper-Level Ontology", in C. Welty and B. Smith (eds.), *Formal Ontology in Information Systems. Proceedings of the Second International Conference (FOIS '01), October 17-19, Ogunquit, Maine*, New York: ACM Press, 34–46. Outlines a framework for realist formal ontology richer than set theory.

Dement, Charles W. Mairet, Charles E., Dewitt, Stephen E. and Slusser, Robert W. 2001 *Mereos: Structure Definition Management Automation (MEREOS Program Final Technical Report),* available through the Defense Technical Information Center (http://www.dtic.mil/), Q3 2001. Example of Ontek work in public domain.

Dipert, Randall R. 1993 *Artefacts, Art Works and Agency,* Philadelphia: Temple University Press. Defends a view of artefacts as products of human intentions.

Forguson, Lynd 1989 *Common Sense*, London/New York: Routledge. Survey of recent work on common sense as a universal of human development and on the relevance of this work for the philosophy of common-sense realism.

Genesereth, Michael R. and Nilsson, L. 1987 *Logical Foundation of Artificial Intelligence*, Los Altos, California: Morgan Kaufmann. On logic in AI. Contains definition of ‘conceptualisation’ in extensionalist terms.

Gruber, T. R. 1993 “A Translation Approach to Portable Ontology Specifications”, *Knowledge Acquisition*, 5, 199–220. Account of the language *ontolingua* as an attempt to solve the portability problem for ontologies.

Gruber, T. R. 1995 “Toward Principles for the Design of Ontologies Used for Knowledge Sharing”, *International Journal of Human and Computer Studies*, 43(5/6), 907–928. An outline of motivations for the development of ontologies.

Gruber, T. R. n.d. “What is an Ontology?”, http://www-ksl.stanford.edu/kst/what-is-an-ontology.html. Summary statement of Gruber’s defintion of ontology as a specification of a conceptualisation.

Guarino, Nicola 1995 “Formal Ontology, Conceptual Analysis and Knowledge Representation”, *International Journal of Human-Computer Studies*, 43, 625–640. Arguments for the systematic introduction of formal ontological principles in the current practice of knowledge engineering.

Guarino, Nicola (ed.) 1998 *Formal Ontology in Information Systems*, Amsterdam, Berlin, Oxford: IOS Press. Tokyo, Washington, DC: IOS Press (Frontiers in Artificial Intelligence and Applications), 1998. Influential collection.

Guarino Nicola 1999 “The Role of Identity Conditions in Ontology Design”, in C. Freksa and D. M. Mark (eds.), *Spatial Information Theory: Cognitive and Computational Foundations of Geographic Information Science*, Berlin/New York: Springer Verlag, 221–234. On constraints on ontologies.

Guarino, N. and Welty, C. 2000 “A Formal Ontology of Properties”, in R. Dieng and O. Corby (eds.), *Knowledge Engineering and Knowledge Management: Methods, Models and Tools. 12th International Conference* (EKAW 2000), Berlin/New York: Springer: 97–112. A common problem of ontologies is that their taxonomic structure is poor and confusing as a result of the unrestrained use of subsumption to accomplish a variety of tasks. The paper provides a framework for solving this problem.

Hayes, Patrick J. 1985 “The Second Naive Physics Manifesto”, in J. R. Hobbs, and R. C. Moore (eds.) *Formal Theories of the Common-Sense World*, Norwood: Ablex. 1–36. Against toy worlds methods in AI, in favour of massive axiomatisation of common sense physics along lines similar to much current work in ontology.

Ingarden, Roman 1964 *Time and Modes of Being*, translated by Helen R. Michejda, Springfield, Ill.: Charles Thomas. Translated extracts from a masterly four-volume work in realist ontology entitled *The Problem of the Existence of the World* (the original published in Polish and in German).

Ingarden, Roman 1973 *The Literary Work of Art: An Investigation on the Borderlines of Ontology, Logic, and Theory of Literature,* Evanston: Northwestern University Press, 1973. Ontology applied to cultural objects.

Johansson, Ingvar 1989 *Ontological Investigations. An Inquiry into the Categories of Nature, Man and Society*, New York and London: Routledge. Wide-ranging study of realist ontology.

Johansson, Ingvar 1998 “Pattern as an Ontological Category”, in Guarino (ed.), 86–94. Ontology of patterns.

Keil, Frank 1979 *Semantic and Conceptual Development: An Ontological Perspective*, Cambridge, MA: Harvard University Press. Study of cognitive development of category knowledge in children.

Koepsell, David R. 2000 *The Ontology of Cyberspace: Law, Philosophy, and the Future of Intellectual Property,* Chicago: Open Court. A contribution to applied legal ontology with special reference to the opposition between patent and copyright.

Koepsell, David R. (ed.) 1999 *Proceedings of the Buffalo Symposium on Applied Ontology in the Social Sciences* (*The American Journal of Economics and Sociology*, 58: 2). Includes studies of geographic ontology, the ontology of economic objects, the ontology of commercial brands, the ontology of real estate, and the ontology of television.

Mertz, D. W. 1996 *Moderate Realism and Its Logic*, New Haven, CN: Yale University Press. Study of the logic and ontology of instantiation.

Mulligan, Kevin 1987 “Promisings and Other Social Acts: Their Constituents and Structure”, in Kevin Mulligan (ed.) *Speech Act and Sachverhalt. Reinach and the Foundations of Realist Phenomenology*, Dordrecht/Boston/Lancaster: D. Reidel, 29–90. On the ontology of speech acts as a foundation for a general ontology of social institutions.

Omnès, Roland 1999 *Understanding Quantum Mechanics*, Princeton: Princeton University Press. Introduction to the consistent histories interpretation of quantum mechanics.

Peirce, C. S. 1933 *Collected Papers*, Cambridge, Mass.: Harvard University Press.

Quine, W. V. O. 1953 "On What There Is", as reprinted in *From a Logical Point of View*, New York: Harper & Row. Defends a view of ontology as the study of the ontological commitments of natural science.

Rossi Mori, A., Gangemi, A., Steve, G., Consorti, F., Galeazzi, E. 1997 “An Ontological Analysis of Surgical Deeds”, in, C. Garbay, et al. (eds), *Proceedings of Artificial Intelligence in Europe* (AIME ’97), Berlin: Springer Verlag.

Searle, John R. 1995 *The Construction of Social Reality*, New York: Free Press, 1995. An ontology of social reality as the product of collective intentionality.

Simons, Peter M. 1987 *Parts. An Essay in Ontology*, Oxford: Clarendon Press. Logical and philosophical study of mereology.

Simons, Peter M. and Dement, Charles W. 1996 “Aspects of the Mereology of Artefacts”, in: Roberto Poli and Peter Simons, ed., *Formal Ontology*. Dordrecht: Kluwer, 1996, 255–276.

Smith, Barry 2001 “Fiat Objects”, *Topoi*, 20: 2. On fiat and bona fide boundaries, with illustrations in geography and other domains.

Smith, Barry and Brogaard, Berit 2002 “Quantum Mereotopology”, *Annals of Mathematics and Artificial Intelligience*, forthcoming. A theory of the relations between partitions at different levels of granularity.

Smith, Barry and David M. Mark 2001 “Geographical Categories: An Ontological Investigation”, *International Journal of Geographic Information Science*, 15: 7. Study of naïve subjects’ conceptualizations of the geographical realm.

Spelke, Elizabeth S. 1990 “Principles of Object Perception”, *Cognitive Science*, 14, 29–56. On the role of objects (as the cohesive bounded foci of local action) in human cognitive development.

Further Reading

http://www.formalontology.it

Hafner, Carole D. and Natalya, Fridman Noy 1997 “The State of the Art in Ontology Design: A Survey and Comparative Review”, *AI Magazine*, Fall 1997, 53–74.

Mulligan, Kevin (ed.), *Language, Truth and Ontology* (Philosophical Studies Series), Dordrecht/Boston/London: Kluwer, 1992.

Smith, Barry (ed.), *Parts and Moments. Studies in Logic and Formal Ontology*, Munich: Philosophia, 1982.

Christopher Welty and Barry Smith (eds.), *Formal Ontology in Information Systems. Proceedings of the Second International Conference (FOIS '01), October 17-19, Ogunquit, Maine*, New

York: Association of Computing Machinery Press.

Glossary

Adequatist Ontology: a taxonomy of the entities in reality which accepts entities at all levels of aggregation, from the microphysical to the cosmological, and including also the mesocosmos of human-scale entities in between (contrasted with various forms of reductionism in philosophy).

Conceptualization: an abstract, simplified view of some domain that we wish to represent for some purpose.

Domain ontology: the extension or specification of a top-level ontology with axioms and definitions pertaining to the objects in some given domain.

Information systems ontology: A concise and unambiguous description of principal, relevant entities of an application domain. A dictionary of terms formulated in a canonical syntax and with commonly accepted definitions, of such a sort that it can yield a shared framework of knowledge-representation on the part of different information systems communities.

Mereology: the formal theory of part-whole relations, sometimes used as an alternative to set theory as a framework of formal ontology.
Metaphysics: commonly used as a synonym of 'ontology'. Sometimes used to refer to the study of competing ontologies with the goal of establishing which of these ontologies is true of reality.

Ontological commitment: the ontological commitment of a theory (or individual or culture) consists in the objects or types of objects the theory (or individual or culture) assumes to exist.

Ontological engineering: the branch of information systems devoted to the building of information systems ontologies.

Philosophical ontology: a highly general theory of the types of entities in reality and of their relations to each other.

Top-level ontology: the general (domain-independent) core of an information systems ontology.