

Foundations of an ontology of philosophy

Pierre Grenon
Department of Philosophy, University of Geneva

Barry Smith
Department of Philosophy, University at Buffalo

Preprint version of paper in *Synthese*, 2011, 182 (2), 185-204 (Ontology Issue)

Abstract. We describe an ontology of philosophy that is designed to help navigation through philosophical literature, including literature in the form of encyclopedia articles and textbooks and in both printed and digital forms. The ontology is designed also to serve integration and structuring of philosophical literature, and in the long term also to support reasoning about the provenance and contents of such literature, by providing a representation of the philosophical domain that is orientated around what philosophical literature is about.

Keywords: Ontology, Philosophy

We take philosophy to be a field of human activity which leads to the creation of entities of a certain special kind: philosophical entities, such as concepts, theories, doctrines, and methodologies. For our purposes here, what makes these entities philosophical is the fact that they are results or outcomes of philosophical activity. What makes such activity philosophical is something which, for our present purposes, can be seen as being primitive and thus undefined. Thus, we will not enter the debate as to what distinguishes philosophical entities from other entities of similar kinds (for example scientific ones). We merely assume that philosophical activity defines a domain in which we find philosophical entities, and we devote our attention to the question of what kinds of philosophical entities there are and how they are interrelated.

The development of ontologies on the part of computer based knowledge system engineers has become common practice. The results

of their work are used as the basis of controlled vocabularies for the annotation of data and information in very many fields, They serve to make this data more easily retrievable, combinable, and susceptible to automatic reasoning. In what follows we apply analogous techniques to the domain of philosophy.

Philosophical creations are entities of the sort that are documented publicly in philosophical literature, and they are themselves subject to further philosophizing. Philosophy itself however is not the sum total of philosophical writings. Rather it is the process which leads *inter alia* to the creation of such writings. Hence an ontology of philosophy is neither merely nor even primarily a theory of philosophical language or terminology. Rather, an ontology of philosophy is a theory of the kinds of entities found in the philosophical domain and of their interrelations.

The distinctive feature of PhilO, the ontology we present in this paper comes from the methodology used to obtain it. Many ontologies in the field of information science are obtained from the semi-automatic application of natural language processing techniques to large corpora of texts. PhilO, in contrast, is itself the product of a philosophical methodology. The result is, to be sure, rather humble as a work of meta-philosophy. This has to do with a number of methodological principles which we will explain in due course. In particular, it is not to be seen as the product of any fixed doctrine. It is, instead, merely a suggested starting point for what we anticipate will be an arduous long-term endeavor. It is to be viewed also as being in every respect revisable. The creation of ontology artifacts to save retrieval and processing of data is an infant discipline, and in this, as in other domains, we are still learning how best to proceed.

All ontologies in non-trivial domains remain forever works-in-progress, and this is true, too, of the PhilO ontology. What we present here is only a portion of a complete ontology of philosophy (more precisely: it is an ontology that covers the entire domain of philosophy but only in first approximation and only at a general level). Our aims are: i) to present a methodology for building ontologies that inspired by a certain philosophical method (which we believe is generalizable to ontologies of other types and in other domains); and ii) to identify the questions which would need to be addressed in order to further enhance the ontology presented.

1 Why an ontology of philosophy?

The products of philosophical activity are nowadays contained in publications, books, articles, and collections thereof (to some degree

also in videos). They are contained also in documents whose purpose is to summarize, such as textbooks, dictionaries, encyclopedias and collections of abstracts. Increasingly, problems are caused – in this as in other domains – by the fact that there is a large and growing mass of documents and other material which one needs to sift through in order to find philosophical contributions of given sorts.

Currently bibliographical databases such as the *Philosopher's Index* are being used as aids to help in organizing and structuring such resources in such way as to make them more easily navigable.. This includes a list of subject terms used to describe or annotate bibliographical entries, modeled on lists such as the Library of Congress Subject Headings created by librarians. The more a search in the database can rely on such lists of keywords, the greater its likelihood of being successful; this is the rationale behind using such lists. But there is a significant shortfall where searches cannot be performed on the basis of matching strings identical to those which appear in the lists of keywords. The same applies to searching for information in a printed volume by using an index, and it applies even when using online resources such as the *Stanford Encyclopedia of Philosophy*, for example through its table of contents, which impose no control over the terminologies used by the authors of the separate entries. Increasingly, the keywords used in browsing through and more or less efficiently retrieving content from such resources are being compiled into so-called *controlled vocabularies* (controlled by the editors of the corresponding resource on behalf of the relevant disciplinary community). There are two major types of such vocabularies in the philosophical domain:

Unstructured thesauri, which are lists of terms with a low degree of informal organization. For example [Broughton1998] consists primarily of two lists; a list of names of persons frequently mentioned in the *Philosopher's Index* (for example Aristotle, Leśniewski and Spinoza), and a list of so-called 'descriptors', which are terms encountered in the database of bibliographical entries (e.g. 'About', 'Abstract', 'Cigarettes', 'Entailment', 'Fictionalism', 'Moral proof', 'New Zealand', and so on).

Structured thesauri, which are lists of terms with some organization, primarily of a hierarchical sort. For example [Berman2001], which is based on the Library of Congress classification, is similar in content to [Broughton1998], but differs in that its terms are organized into families and ordered (into 'narrower' and 'broader' terms) according to level of generality.

The notion of generality involved here is however still somewhat idiosyncratic, and defies logical definition. Thus, for example, ‘Beauty’ is seen as being a narrower term than ‘Aesthetics’. This does not mean that *beauty* is a subkind or instance of *aesthetics*; rather it means that documents dealing with the concept *beauty* are intended to be included by the compilers of this resource among the documents dealing with *aesthetics*.

While unstructured thesauri are very useful for example in indicating coverage of bibliographic resources via enumeration, they do not convey any further information pertaining to the meanings of the terms they list. Moreover, they typically contain large numbers of terms which do not seem properly to belong to the domain in question. Thus although there may be a number of philosophical publications addressing issues related to cigarettes and smoking, it is unclear whether representations of these items ought to belong to an ontology of philosophy more strictly conceived.

Structured thesauri do carry some further information, in particular they loosely indicate certain forms of relatedness between terms that are not hierarchical (as when saying that the term ‘philosopher’ is a term related to the term ‘philosophy’). It is one shortcoming of these structures that they do not specify further the non-hierarchical relationship between their terms but, again, there is a more fundamental concern which derives from the origin of the mentioned resources in the realm of library science. For as will by now be clear, the information they contain pertains not to the meanings of the included terms, but concerns rather the documents which these terms are used to index. The relation captured in the subordination of Beauty to Aesthetics is something along the lines of: *beauty* is a concept used in works in the philosophical field of *aesthetics*. Thesauri are blind to the structural relations that obtain between the referents of the terms they list, but it is precisely this sort of structure that an ontology of philosophy of the sort we are constructing is in the business of providing.

2 A Philosophical Approach to Ontology Building

Ontologies as information artifacts are constructed nowadays in very many disciplines [Watson n.d.], and methodologies differ as to the sources used and the role of human intervention. We have referred already to the distinction, amongst ontologies in information science, between those that are handcrafted, and those generated via natural

language processing techniques. The latter are in practice created semi-automatically, since the process of ontology extraction requires the validation by human editors if it is to yield usable content. The most successful approach to the building of ontologies seems however still to be one which relies entirely on human input. This is so, for example, of the Gene Ontology and of the other biomedical ontologies now being heavily applied in clinical and translational research [Rubin2008]. Increasingly, the latter are relying on an approach rooted in part in the acceptance of the need to take seriously insights of logicians and philosophers for example on the role and nature of definitions and on issues of meaning and reference [Smith2003, Smith et al., 2007]].

There is a simple rationale for using a philosophical approach to ontology elaboration in whatever one's chosen domain. It is that, through careful examination and logical analysis, and careful attention to potential ambiguities and to the category mistakes and mistakes of use and mention that have plagued ontology construction in many information science circles thus far, we can reach more accurate and consistent representations of the domain at issue and of the relations which obtain between the represented entities, of a sort which is more readily able to support logical reasoning. This same motivation speaks also in favor of ontologies created manually from the start, not least because reviewing the product of automated language processing is a task which, in our experience, rarely leads to outcomes which are structurally sound.

Automated techniques yield networks of 'associated' terms which are thought to be more or less closely related; they yield what are called 'lexical networks'. But such artifacts are no more insightful when it comes to representing the structure of a domain than are the sort of thesauri which relate terms according to their putative co-occurrence in an indexed document. Terminological and lexical information based on co-occurrence links are useful for certain retrieval purposes, but they do not provide a reliable representation of the corresponding target domain, and they do not provide an account of how entities in that domain are interrelated.

One further problem pointing to the limitations of lexical approaches is the lack of interoperability. This is because, even where one and the same term appears in a plurality of such systems, there is no guarantee that it will be similarly handled. One important quality criterion on ontologies, however, is that ontologies should as far as possible embrace a principle of orthogonality (meaning: convergence on a

single ontology for each domain) and that ontologies for neighboring domains should work well together [Smith2008]. The philosophical approach we advocate rests on a view of ontologies as consisting of representations of the entities in the domain of reality to which the ontology relates. Only on the basis of representations of this sort, we believe, will it be possible to make coherent progress in linking together different terminology systems (for example in different languages).

3 Guiding principles

Our methodological approach is perhaps best summarized by a number of guiding principles for ontology building.

Realism

Ontology, as we conceive it, is concerned with providing an account of the entities existing within a given domain of reality, where ‘reality’ is here understood in the broadest possible sense, to include for example not only molecules and planets but also works of literature, laws, and historical epochs. The objects of the ontological inquiry into a domain D are first-order entities in the domain D, rather than concepts in the minds of people (experts, in particular) who study D or terms used (by experts, in particular) to refer to D and its components.

Concepts and terms may, though, perfectly well form the subject matter of ontologies addressing psychological or linguistic domains; then, however, they are first-order entities in their own right. In the domain of philosophy, of course, many entities are concepts which our ontology is intended to help categorize.

Relevance and modularity

Before we can embark on the construction of an ontology of philosophy, we need to establish what sorts of entities and relations exist in the philosophical domain. This is problematic in part because many of these entities fall under kinds which are in fact contextual specializations of more generic kinds, and pinpointing the differentia for the more specialized kinds is one important part of the ontology enterprise. Thus for example the kind *philosophical concept* and the kind *philosophical theory* are *prima facie* formed respectively by just those concepts and theories which are philosophical. What this means, however, is far from being trivial.

Establishing what sorts of entities and relations exist in the

philosophical domain is problematic also because there are entities that may not be specific to the domain of philosophy but appear only under a certain guise in this domain. For example Bertrand Russell was a philosopher at certain intervals in his life, but he was not born a philosopher. He was also a father, an Englishman and many other things that are beyond the purview of an ontology of philosophy. Philosophers are all those persons who are involved in some way in the domain of philosophy. But they do not form a natural kind. To be a philosopher is what is sometimes called a role and typically demands a relational account (e.g. in terms of the participation of role-bearers in certain activities) [Trautwein and Grenon, 2003; Arp and Smith, 2008]. To ease our problems with such questions we adopt two fundamental principles:

The *principle of relevance*: we are interested in entities or features of entities which belong exclusively to our selected domain. For example we are interested in Bertrand Russell's philosophical activity and productions and not in his biography as a political activist. Also we are interested in philosophical concepts, not in concepts as such.

The *principle of modularity*: we assume that our ontology of philosophy is to be integrated into a larger body of interoperable ontologies pertaining to other, neighboring domains, for example , culture, politics, science, history, literature, theology. It is in this larger embedding system that categories such as person, for example, would be found, thereby enabling us to attach to Bertrand Russell his personal features. This allows us also to make provision for fitting our ontology of philosophy under a more general umbrella ontology in which the generic features of concepts could be accounted for and in which also distinctions such as that between concepts and theories could be made in a more robust fashion.

While the principle of relevance is used to select elements to include in the ontology of philosophy, the principle of modularity is there to allow room for elements that will allow us, in the future, to complete and embed the representations in this ontology into a broader system.

Maximally opportunistic use of resources

In the main, our method is to proceed from ground-level analysis of the alleged entities in a given domain (for example, the *philosopher* Bertrand Russell, the *concept* of definite description or the *axiom* of reducibility) to the elaboration of a system of kinds of entities and their

relations. The question we face now is: which sources and resources we should use for this purpose. This is not a simple question, because resources may differ not only in quality and comprehensiveness, but also in the sorts of biases they impose (for example Western vs. non-Western, analytical vs. Continental, and so forth).

For the purpose of initial term selection there is some value in artifacts such as thesauri. The task of sketching an ontology can partly be seen as one of sifting through and organizing the lists of philosophical entities which such thesauri, in their ramshackle way, represent, into coherent categories organized hierarchically by type and subtype. This is not, however, a fully satisfactory strategy because such lists fail to account for the nature of philosophy as a complex domain, in which the different sorts of entities are related together by ontologically important sorts of relations (for example of parthood, precedence, influence).

We may also draw on sources such as textbooks and articles for term selection. Unfortunately these, too, differ in the way they recognize alleged entities in the domain of philosophy and in the way they partition the domain of philosophy itself. Here, moreover, there is the problem of factual accuracy and also doctrinal neutrality, so that the question arises as to what sources can be trusted and to what degree.

These considerations suggest a combination of an empirical approach, starting from a variety of established lists of relevant entities in the form of abstracts repositories and textbook indices in the domain of philosophy, supplemented by a more global classificatory approach to the domain – based on logical principles that are as far as possible neutral as between different points of view – the latter to be used as a means of ensuring consistency and coherence of the ontology structure.

Philosophical neutrality

Initially we rely on those sources that are commonly recognized as authoritative. In many domains, such as those of the sciences or engineering, ontologists consult experts from whom they elicit knowledge about the domain. One could argue that this procedure is compromised in our present case, given the conceptual and controversial character of philosophy. But an ontology of philosophy does not have to engage with or resolve philosophical disputes. Rather, it is concerned with what entities there are in the domain of philosophy and thus also with what entities philosophical debates are concerned. Thus an ontology of philosophy has to be guided by a principle of neutrality regarding its content in order to make room for all

philosophical views, the latter themselves, together with the associated disputes, being treated as entities in their own right. The driving force behind the adoption of an unbiased perspective on philosophy is at bottom to ensure the adequacy of our representation.

Revisability of the representation

One of the problematic aspects of many engineered ontologies is that they are static artifacts. This has to do with the fact that they need to be used in robust software applications to meet well-defined goals. They are not easily modified nor amended, because they are tailored to be used by some specific system. We can see an analogous phenomenon also at the level of library artifacts such as thesauri and classification systems. The list of subject headings used by the Library of Congress, for example, has remained fundamentally unchanged for more than a century, and thus still shows considerable influence of 19th century scholarship in the US. Problems arising from such legacy phenomena give rise to short term *ad hoc* solutions when unforeseen situations are encountered (for example the need to classify books on hitherto topics). As ontologies have become increasingly adopted by systems requiring frequent updates, ways have had to be found to design them in such a way that they are easily extended and revised.

The reasons to allow for ontology change turn not only on the fact that our knowledge is growing and being constantly subjected to correction, but also on the fact that the world is changing. The changes affect not only the world of information artifacts which some ontology terms will be used to describe, but also the world that is represented in these artifacts. Ontologies rest on accounts of reality which are based on expert knowledge, and not only knowledge of reality can evolve but so also can reality itself. This is true, too, in a domain such as philosophy.

One objection which may be made against our approach is that the needed philosophical neutrality of the representation is betrayed by our adoption of what amounts, in effect, to a principle of realism as concerns the entities that populate the domain of the ontology we are constructing. Perhaps it is a philosophically biased position which regards (philosophical) concepts, theories, disputes, and so forth as entities. But this is the sort of bias that comes with the territory. Ontologies *are* representations of entities and of the relations between them. What we should beware of doing is building into our ontology controversial relations between the concepts which the ontology treats (for example that beauty is a kind of good). On the other hand, where

claims are controversial, we must have the possibility of representing the distinction between doctrines which embrace these claims and doctrines which deny them.

4 The big picture

The first thing we need to do in building an ontology of philosophy is to delimitate the domain. The most generic claim behind our choice is that philosophy is an activity carried out by human beings, and that the main output of this activity is entities of certain sorts: philosophical entities. We believe that the thesis that philosophy is an activity carried out by human beings is uncontroversial. But even if we are wrong, a combination of our principles of modularity and revisability will allow a broader category of entities to be recognized as the substrate for the role of philosopher (for example including non-human agents as authors of philosophical ideas).

Another potential problem has to do with how we view the outputs of the activity of philosophizing. For example, is producing a philosophical concept a case of creation or discovery? We leave open such questions by employing a correspondingly general reading of ‘output’. concepts, theories, arguments and methods

The approach we have sketched so far leads us to identify three main features characterizing the domain of philosophy. In the first place there are philosophical entities which may belong to various categories (for example of concepts, theories, arguments, methods). In the second place there is philosophy itself, which is an entity dependent on certain activities performed by philosophers. We can then easily see, that philosophy is a field dividing into subfields (e.g. metaphysics, philosophy of science, aesthetics, and so forth). Finally, we can see that there are philosophers. We have said already that philosophers do not form a natural kind, but it is nonetheless also not only possible but traditional to divide philosophers into various groups according to more or less robust criteria (e.g. community of thought, tradition, nationality or period of principal activity).

For the purpose of representing the domain of philosophy and philosophical activity in the form of an ontology, we are concerned not with the nature of philosophy as such, but rather with the distinctions and interrelations between the categories identified thus far. It will matter to us, when mapping (as in: drawing a map of) the philosophical domain that there are at least three main polarities for such an ontology – philosophers, philosophical entities, and the field that is covered by

these entities (what these entities are about). But we can leave open also the question of their precise nature. Whatever the answers to these questions might be, our claim is that the distinctions and interrelations between these three groups of entities will be preserved (no field is a philosopher, no philosopher is a philosophical entity, and no philosophical entity is a field).

4.1 Philosophical entities

Philosophical entities are those entities which live their life within the philosophical domain of reality. They do not come in one kind only. Some are more simple, some more complex, and we can distinguish part-whole relations between philosophical entities of certain sorts, as for example between the concept of space and a theory of space. Such relations will themselves be of different kinds, so that we have not only different degrees of complexity but also different kinds of complexity.

A cursory survey of philosophical activities suggests the following preliminary list of kinds of philosophical entities:

- concept (e.g. the concept of form)
- proposition (e.g. that forms exist)
- theory (e.g. Plato's theory of forms)
- argument (e.g. Plato's Third Man argument)
- method (e.g. the dialectic method).

Philosophical concepts are as near as we shall come to basic units of philosophical activities. Philosophical propositions are in first approximation made of concepts. And theories are made of propositions. There are also a number of roles that some of these objects may take on, for example, with respect to argumentation, the role of an *axiom*, *hypothesis*, *theorem* or *conclusion*. Arguments are of course primarily logical entities, but they can be philosophical in two ways: firstly when their object is philosophical and secondly when they follow argumentation patterns which are properly philosophical. This may be the case when arguments contain propositions whose support is philosophical in nature (e.g. appeals to intuition). Methods, too, are not in and of themselves philosophical, but there are philosophical methods (e.g. introspection, phenomenological analysis, argumentation, conceptual analysis, and so on), which will fall within the coverage domain of philosophical ontology as here conceived.

4.2 Philosophical fields and philosophers

Philosophy, the field of activity, is a particular entity. It can be broken down into branches, its subfields. There is also a way of partitioning the field of philosophy in order to do justice to the fact that there is philosophy that is the philosophy *of* some philosopher. We make room for this partition in two ways. Consider young Immanuel, mesmerized by the beauty of a pebble in the garden, and wondering whether the hidden face of the pebble exists. Immanuel is here philosophizing, though he does not yet know that he will grow to be the great philosopher Immanuel Kant, who will be remembered for only a tiny part of his lifelong philosophizing activity. He does not suspect, either, that out of this already impoverished portion of this whole, some subportion will be more or less digested by generations of philosophers to come, who will produce the philosophical entity called ‘Kantian philosophy’. There is Kant’s philosophizing, his philosophy, there is some congeries of theories that is an output of this activity, and there are various sequelae of this output.

As philosophy is also occupied with itself, there are two specific kinds of reflective exercise concerned with Kant’s philosophy. There is, on the one hand, philosophical historiography concerned with the philosophizing activity of Kant himself, from pebble to death. And on the other hand, there is Kantian philosophy, a developing creature with a life of its own. The first is concerned with a subfield of philosophy in which Kant himself was the main, if not the only player; the second is concerned with a theory or family of theories, and thus with philosophical entities in our sense.

Entities like the philosophy of Kant (in the first sense distinguished above) are not prime examples of philosophical fields; they are in fact very special cases. More interesting to the ontology of philosophy for us here are such portions of philosophy as are concerned with knowledge (epistemology) or with science (philosophy of science) or even with Kant’s writings (Kantian exegesis). These are distinctions among philosophical fields according to the topic with which the philosophizing activity is concerned. Such topical distinctions form the more robust and bona fide distinctions among subfields of philosophy.

It is traditional to speak of philosophy as having a variety of branches or subdisciplines. Although there is a handful of examples of a very generic nature on which everybody seems to agree (metaphysics, ethics, political philosophy, philosophy of science, and so on), it is a quite obscure matter to recognize how philosophy is divided and by

what principles. In section 6 we will propose systematic criteria for dividing philosophy into subfields. We will also see how some of these criteria can be applied in the classification of philosophical entities of other sorts, including philosophers themselves. Classifying philosophers is warranted not only because philosophers are sources of philosophical entities – and thus of philosophy and its subfields – but also because their works, as well as to some extent they themselves, are subjects of further philosophizing. Thus there are not only ethicists and metaphysicians, but also Aristotle scholars and Hegel exegetes as well as philosophers of sport and philosophers of engineering. This does not mean that philosophers form kinds; rather they enter into, or form, groups. Thus Kant is not an instance of a putative philosopher-kind, not least because he wasn't born a philosopher, however early he might have started philosophizing. Kant, like Plato, belongs to the group of people who are philosophers. This in turn means that he engaged in philosophical activity. Philosophical fields are niches for philosophers to produce philosophical entities, and these entities in turn may serve as tools (inputs) for further philosophizing.

Figure 1 sums up in visual form the big picture that emerges from the foregoing discussion. The main kinds of entities we are interested in the philosophical domain are branches of philosophy, philosophers and groups thereof, in so far as they bring about and deal with philosophical entities, and these philosophical entities themselves.

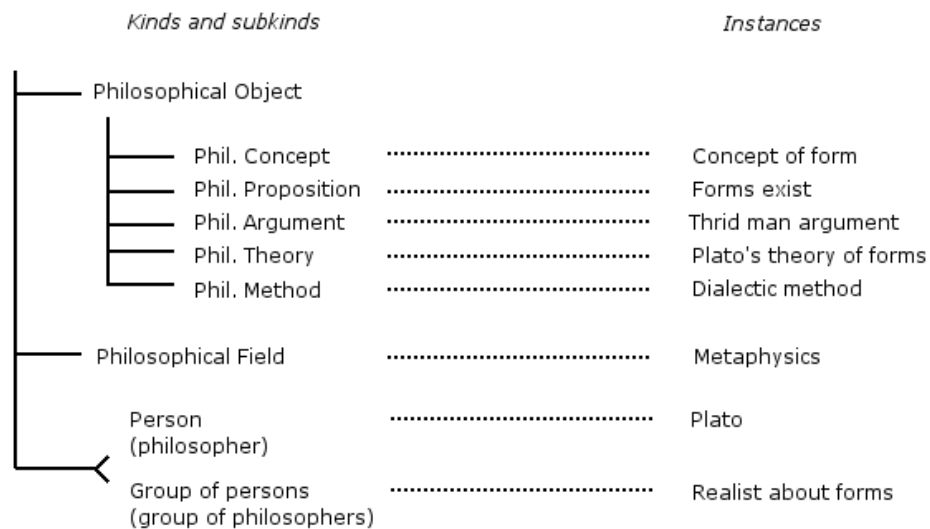


Figure 1: Main kinds of entities in the domain of philosophy and examples of their instances.

5 Classificatory distinctions

Subfields of philosophy, philosophers and groups thereof and philosophical entities (concepts, theories, and so on) are the main kinds of entities we find in the philosophical domain. But of course it is possible, and indeed traditional, not only to arrange philosophy into subfields, but also to arrange philosophers into subgroups and, finally, to divide philosophical entities, concepts in particular, into a variety of kinds. Thus philosophy is divided into branches addressing specific topics, following specific methods or adhering to specific perspectives. Philosophers are often classified on the basis of the historical period during which they lived but also of the country in which they were born or in which they were active as philosophers and the languages in which they spoke or wrote. Philosophical entities are in turn attached to these divisions, as the concept of *nous* is an example of an ancient Greek philosophical concept and the concept of *transcendental ego* is an example of German idealist concept.

These divisions do not however yield clear cut and systematic distinctions. In order to systematize the divisions in question we appeal to our principle of modularity. There is, in the first place, a backbone classification which corresponds to the main categorial distinctions we have made in the foregoing (into concepts, theories, and so forth), together with a subdivision of philosophy itself which follows the traditional topic-oriented division (into ethics, aesthetics, and so forth). In addition, we make provision for a number of modules providing further dimensions onto which to project this backbone structure. Suppose for example that we have an ordered list of periods of time. It is deceptively simple to segment philosophy as a whole, and many of its subfields, and many groups of philosophers, to produce a corresponding set of time-determined divisions: 20th-century ethics, 19th-century Aristotelians, and so forth. And what can be done with time can be done with many other aspects of philosophical activity and its actors, including geographical location, nationality, and cultural, linguistic or religious background.

Although a formal approach allows systematization of such divisions along combinatoric lines, it is unclear to what extent the results are coherent reflections of genuine divisions in the underlying domain. Consider for example what German philosophy is or what a German philosopher is. Are these terms referring to a geographical, a geopolitical or a linguistic division? When does Germany start and end historically? Is Kant a Prussian or a German philosopher, and is he

somehow less German than Heidegger or Habermas? And if the feature in question is linguistic, then is Roman Ingarden a German philosopher by this light?

Divisions are easily defined combinatorically. The challenge, however, is to see whether such divisions are actually significant for carving up the domain at its joints, and this is hard work, as is shown for example by the discussion of the meaning of ‘Austrian philosophy’ in [Simons2004]. It would be the work of many disciplines perhaps (sociology, history, philosophy itself) to evaluate these divisions as perspicuous and tenable. It is in any event an enormous enterprise (see for example [Holenstein2004]). What we are interested in here is the power of formal ontology in generating and manipulating such divisions. The question of which of them has a counterpart in reality is an empirical matter and hence a question that has to be answered by other means.

6 Formalization

PhilO is a formal ontology of the philosophical domain. Formalization is useful on many counts. Not least because it is one step towards making the product of an ontological investigation readily available to computer processing, but also, for explanatory purposes in virtue of the clarity of expression it provides. For our purposes here we adopt the resources of classical first-order predicate logic (with identity) and use the usual symbols for logical constants (in particular, ‘ \rightarrow ’ for material implication). We use concatenations of large and small capital letters for names of individual constants (e.g. PHILOSOPHY but also PHILOSOPHICALFIELD) and the letters of the end of the Latin alphabet for variables ranging over individuals. Predicates will be italicized strings of letters of the alphabet (e.g. *instantiates* but also *workedOn*). Finally, we omit external universal quantifiers.

6.1 Top-level categories in PhilO

An ontology is first and foremost a theory of entities, their kinds and their interrelations. We need now to put in place the formalism allowing us to sketch such a theory for philosophical entities.

Instantiation of a kind and subsumption between kinds

The first relation that occurs in an ontology is that between a kind and the entities that fall under this kind. This is the relation of instantiation (here dubbed *instanceOf*) holding, for example, between philosophy

and the kind philosophical field or between Immanuel Kant and the kind philosopher. We will leave open the question of the features of the relation of instantiation, but we will treat all entities, both particulars and kinds, as individuals in our domain, following the strategy outlined in [Smith2005] and applied to the domain of biology in [Smith *et al.* 2005].

An ontology arranges kinds into classificatory trees or ‘taxonomies’. Taxonomies are ordered by a relation *subcategoryOf* between categories (commonly called ‘*is_a*’ in ontology engineering circles). Thus the second relation that occurs in an ontology is the relation of subsumption between kinds which holds, for example, between the kind philosophical concept and the kind philosophical entity. The relation of subsumption among kinds can be defined in the following way:

We will use *disjoint* for expressing the relation between two or more kinds when they share no instances).

$$\textit{subcategoryOf}(x, y) \equiv_{\textit{def}} \forall x (\textit{instanceOf}(z, x) \rightarrow \textit{instanceOf}(z, y))$$

We can now register what we have said in our informal discussion above:

instanceOf(PHILOSOPHY, PHILOSOPHICALFIELD)
instanceOf(KANT, PHILOSOPHER)
disjoint(PERSON, PHILOSOPHICALFIELD, THEORY)
subcategoryOf(CONCEPT, PHILOSOPHICALENTITY)
subcategoryOf(PROPOSITION, PHILOSOPHICALENTITY)
subcategoryOf(THEORY, PHILOSOPHICALENTITY)
subcategoryOf(ARGUMENT, PHILOSOPHICALENTITY)
subcategoryOf(METHOD, PHILOSOPHICALENTITY)
disjoint(CONCEPT, PROPOSITION, THEORY, ARGUMENT, METHOD)

Not all hierarchical structures in an ontology are subsumption trees. For categories can be also organized, for example, according to how their instances relate through the relation of parthood (paronomies). More generally, the relations there are between instances of kinds will allow for defining various relations between these kinds [Smith *et al.*, 2005]

6.2 Domain relations

We have already alluded to a number of ontological relations in the domain of philosophy. For lack of room, we will only draw a formal characterization of a small number of them. It is already sufficient to go

a long way in representing the domain of philosophy and we hope also that it will indicate the right direction for extending this preliminary account.

	group	concept	field
person	<i>memberOf</i>	<i>workedOn</i>	<i>activeIn</i>
group	<i>subgroupOf</i>	<i>workedOn</i>	<i>activeIn</i>
concept	-	<i>subconceptOf</i>	<i>infield</i>
field	-	-	<i>subfieldOf</i>
philosophical entity	-	-	<i>inField</i>

Table 1: Examples of binary relations used in PhilO. Entries in the first column indicate the *domain*; entries in the first row indicate the *range* of the corresponding relation.

Subsumption among fields

We use *subfieldOf* for the relation between two fields when the first is more specific than the second as, for example, between metaphysics and philosophy. This relation satisfies the following axioms

$$\begin{aligned}
& \text{subfieldOf}(x, y) \rightarrow (\text{instanceOf}(x, \text{PHILOSOPHICALFIELD}) \wedge \\
& \text{instanceOf}(y, \text{PHILOSOPHICALFIELD})) \\
& \neg (\text{subfieldOf}(x, y) \wedge \text{subfieldOf}(y, x)) \\
& (\text{subfieldOf}(x, y) \wedge \text{subfieldOf}(y, z)) \rightarrow \text{subfieldOf}(x, z)
\end{aligned}$$

Subsumption among concepts

We use *subconceptOf* as the relation between two concepts when the first is a specialization of the second. We mean moreover that the specialization is definitional and not subject to philosophical debate, thus, in particular, not theory dependent. For example, the concept of feminine beauty is a sub-concept of that of beauty. But also the concept of space in aesthetics is a subconcept of the concept of space in philosophy *simpliciter*. However, the concept of beauty is not a sub-concept of the concept of good, nor is the concept of person a sub-concept of the concept of material object, irrespectively of whether beauty is good or a kind of good or of whether persons are material objects. *subconceptOf* is a relation between concepts that is asymmetric, and transitive:

$$\text{subconceptOf}(x, y) \rightarrow (\text{instanceOf}(x, \text{PHILOSOPHICALCONCEPT}) \wedge$$

$$\begin{aligned}
& (\text{instanceOf}(x, \text{PHILOSOPHICALCONCEPT}) \\
& \neg (\text{subconceptOf}(x, y) \wedge \text{subconceptOf}(y, x)) \\
& (\text{subconceptOf}(x, y) \wedge \text{subconceptOf}(y, z)) \rightarrow \text{subconceptOf}(x, z)
\end{aligned}$$

Group membership and subsumption among groups

We use *memberOf* for the relation between a person and a group (as between Kant and the group of all philosophers) [Bittner *et al.* 2004]. All groups in the domain of the PhilO ontology will be sub-groups of the group of all philosophers. We use *subgroupOf* for the relation between two groups of philosophers (more generally: between groups of persons) when the first is a group included in the second (the Cynics formed a group of philosophers in the here intended sense). The relation *memberOf* is irreflexive:

$$\neg (\text{memberOf}(x, y) \wedge \text{memberOf}(y, x))$$

We can define *subgroupOf* as follows:

$$\text{subgroupOf}(x, y) \equiv_{\text{def}} \forall z (\text{memberOf}(z, x) \rightarrow \text{memberOf}(z, y))$$

Relation of an entity to a philosophical field

We use *inField* as a generic relation between a philosophical entity and a philosophical field when the entity is one that belongs to that field as, for example, the concept of beauty belongs to the field of aesthetics.

$$\begin{aligned}
& \text{inField}(x, y) \rightarrow (\text{instanceOf}(x, \text{PHILOSOPHICALENTITY}) \\
& \wedge \text{instanceOf}(y, \text{PHILOSOPHICALFIELD}))
\end{aligned}$$

We keep this relation generic for the purpose of this presentation but of course what it is for a philosophical entity to belong to a field comes in many flavors not least because philosophical entities themselves come in different kinds. Compare for example the relation between the concept of beauty and the field of aesthetics to the relation between the concept of space and the field of aesthetics and compare them both to the relation between Kant's theory of aesthetics to the field of aesthetics.

All philosophical entities in a field are entities in any superfield of that field.

$$(\text{subfieldOf}(x, y) \wedge \text{inField}(z, x)) \rightarrow \text{inField}(z, y)$$

Moreover, all philosophical entities are objects in at least one field, even if this be only the field of philosophy.

$$\text{instanceOf}(x, \text{PHILOSOPHICALENTITY}) \rightarrow \exists y \text{inField}(x, y)$$

Activity of a philosopher or group thereof in a field.

We use *activeInField* for the relation between a philosopher and a philosophical field in which the philosopher in question has been or is active as for example, Kant has been active in the field of aesthetics. Use of this relation allows for registering contributions by a person to a philosophical field. It is generic in the sense that it does not specifically relate to the contribution itself. Moreover, being active in a given field transfers to any superfield of that field.

$$(subfieldOf(x, y) \wedge activeInField(z, x)) \rightarrow activeInField(z, y)$$

We can remark that if a person is active in any philosophical field, she is *ipso facto* a philosopher. Generalizing that notion, we can also relate a group of philosophers and a philosophical field of their activity. For example, ethicists are those philosophers active in the field of ethics. We have a variety of options available as to how to represent such relations. For the sake of simplicity, we will allow expressions such as '*activeInField*(ETHICIST, ETHICS)'. In order to have a finer grain representation, one could introduce variants of *activeInField*, one for individual philosophers and one for groups.

Work on philosophical entities

We use *workedOn* as the relation between a philosopher and a philosophical entity. This is a generic relation between person (here: a philosopher) and a philosophical concept. This is for example the relation between Kant and the concept of transcendental ego or Kant and the theory of transcendental aesthetics. Here too we could refine the representation by introducing variants in order to account in particular for the variety of ways in which the *workedOn* relation between a philosopher and some philosophical entity can obtain. It would seem that the most important relation between philosophers and concepts is a relation of creation. Indeed, philosophy as an activity is in large part the creation of concepts. But philosophy is also a public matter and there are thorny issues as to the metaphysical nature of concepts themselves. This poses problems for the introduction of a creation relation, for instance because it is often indeterminate whether a given concept should be described as having been created anew by a given philosopher or rather appropriated or rediscovered. For these and similar reasons we confine ourselves here to documenting the relation of working on.

$$(workedOn(x, y) \wedge objectInField(y, z)) \rightarrow activeInField(x, z)$$
$$activeInField(x, y) \rightarrow \exists z (workedOn(x, z) \wedge objectInField(z, y))$$

6.3 Further classificatory elements

In addition to the categories and relations we find in the domain of philosophy, it is possible to carve out further distinctions for example among entities such as fields, philosophers, and objects. We will here only sketch how this may be done in modular fashion.

Suppose we have a dimension such as time along which we wish to segment a number of philosophical fields. We can then map the one onto the other in the obvious way, creating terms such as ‘20th century philosophy’, ‘19th century metaphysics’, ‘18th century ethics’, and so forth. Now, imagine we do the same with another dimension such as one that lists cultural or national groupings. We can easily combine these in order to produce a more complex segmentation. Such a process can be reiterated indefinitely up to a level at which we reach very specific segments. This process is illustrated in Figure 2.

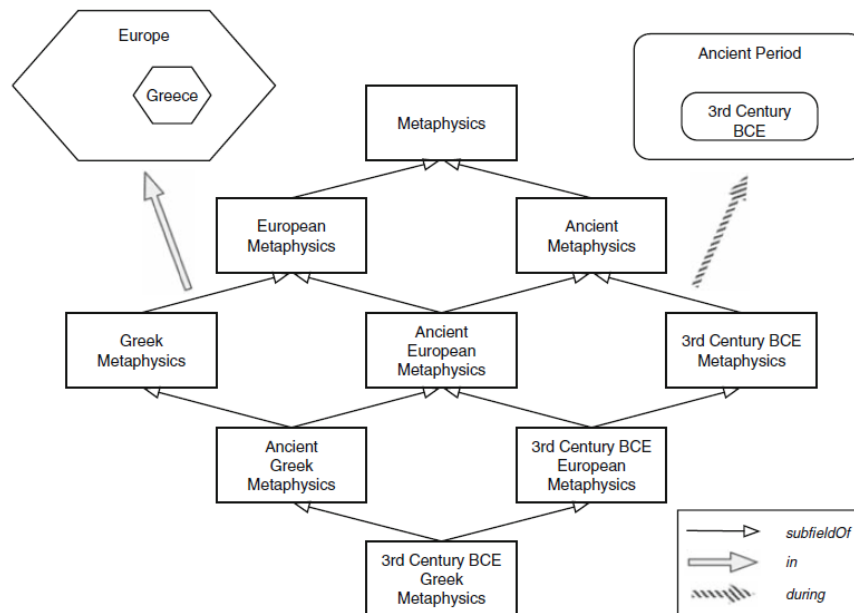


Figure 2: Metaphysics segmented along geographical and temporal axes with the resultant cross-segmentations

As already noted, such systematic segmentation might produce segments which are not all equally interesting from a philosophical standpoint or from the standpoint of history of philosophy. However, deciding to which degree the delineation of a given segment is relevant,

insightful or valuable belongs to a level of analysis more detailed and more empirically orientated than that attempted here. PhilO, ultimately, must absorb this more detailed level. The preliminary version of the ontology presented here is however important, since it allows this more profound kind of segmentation to take place.

Figure 3 provides an overview of the types of information whose representation is supported by PhilO already in its preliminary form presented here.

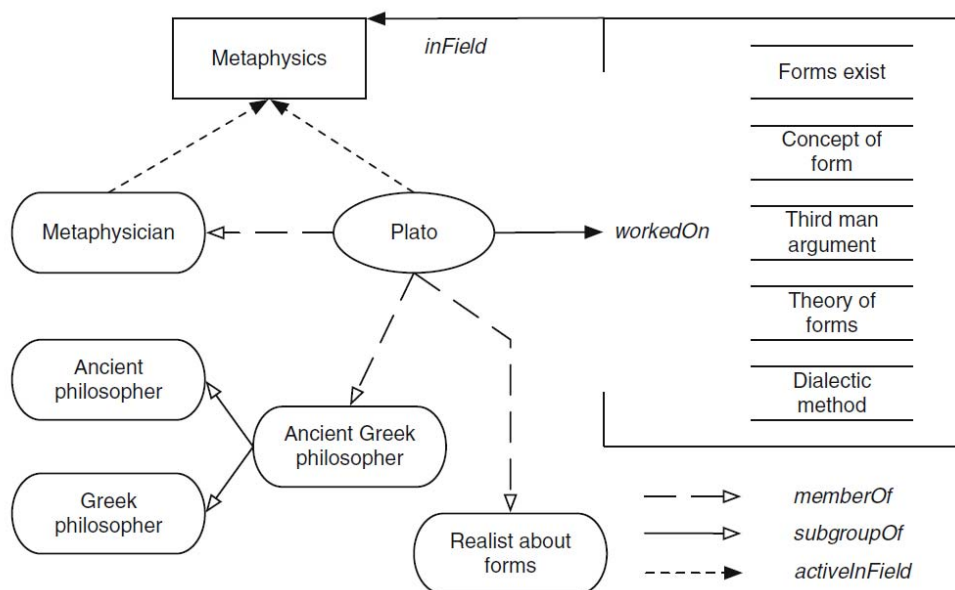


Figure 3. Plato and some of the entities surrounding him in the domain of philosophy.

7 Conclusions and future work

We have presented a preliminary version of the PhilO ontology. In virtue of the methodological approach guiding its elaboration this preliminary version is just a step towards a fuller treatment of the whole domain of philosophy. Because of the modularity of our approach, we can take for granted that much relevant material from neighboring ontologies – for example from the domain of geopolitical ontology, ontology of persons, of activities, of publications and so on – could be added to PhilO by aggregation. One line of future work would

then be to link up the elements presented here to such external modules, eliciting in particular significant cross-ontology relations and adding further axioms to fit. But there is also much work to be done on the PhilO ontology itself. We have indicated in several places simplifications and approximations. In particular, we have only provided a sketch of the formalization of a selected fragment of the relations which show up in the domain among those categories we have singled out. In several cases, we have indicated desirable specializations or generalizations of both categories and relations. Once the catalogue has been augmented there will follow the arduous task of axiomatizing in a more detailed and serious fashion. Finally, the purpose of an ontology is to provide the elements for the representation of a domain and the ultimate test for its validity will be to try registering knowledge on a broad scale using its terms to some set level of detail, and to demonstrate pragmatic benefits from such segmentation of the sort which are already being harvested in other domains.

Acknowledgments

This paper is based on work supported by the Philosophy Documentation Center.

References

- [Arp2008] Arp, R. and Smith, B., “Function, Role, and Disposition in Basic Formal Ontology”, *Proceedings of Bio-Ontologies Workshop* (ISMB 2008), Toronto, 45-48. Available at <http://bio-ontologies.org.uk/download/Bio-Ontologies2008.pdf>.
- [Berman2001] Berman, B. L. (ed.) *Library of Congress Subject Headings in Philosophy: A Thesaurus*. Philosophy Documentation Center, 2001.
- [Bittner et al. 2004] Bittner, T., Donnelly, M. and Smith B., “Individuals, Universals, Collections: On the Foundational Relations of Ontology”, in Achille Varzi and Laure Vieu (eds.), *Formal Ontology and Information Systems*, Amsterdam: IOS Press, 2004, 37–48.
- [Broughton1998] Broughton, K. M. (ed.) *The Philosopher’s Index Thesaurus*. Philosophy Documentation Center, 1998.
- [Holenstein2004] Holenstein, E. *Philosophie–Atlas: Orte und Wege des*

- Denkens* Ammann Verlag, 2004.
- [Rubin2008] Rubin, D. L., Shah N. H. and Noy N. F. “Biomedical Ontologies: A Functional Review Briefings in Bioinformatics”, 9 (2008), 75-90.
- [Simons2004] Simons, P. *The Anglo–Austrian Analytic Axis Philosophy and Logic in Central Europe from Bolzano to Tarski*. Dordrecht: Kluwer, 2004.
- [Smith2003] Smith, B. “Ontology”, in Luciano Floridi (ed.), *Blackwell Guide to the Philosophy of Computing and Information*, Oxford: Blackwell, 2003, 155–166.
- [Smith2005] Smith, B. “Against Fantology”, in Johann C. Marek and Maria E. Reicher (eds.), *Experience and Analysis*, Vienna: HPT&ÖBV, 2005, 153–170.
- [Smith2008] Barry Smith, “Ontology (Science)”, in C. Eschenbach and M. Gruninger (eds.), *Formal Ontology in Information Systems*, Amsterdam: IOS Press, in press.
- [Smith et al., 2007] Smith, B., Ashburner, M., Rosse, C., Bard, J., Bug, W., Ceusters, W., Goldberg, L. J., Eilbeck, K., Ireland, A., Mungall, C. J., The OBI Consortium, Leontis, N., Rocca-Serra, P., Ruttenberg, A., Sansone, S.-A., Scheuermann, R. H., Shah, N., Whetzel, P. L., Lewis, S. “The OBO Foundry: Coordinated Evolution of Ontologies to Support Biomedical Data Integration”, *Nature Biotechnology*, 25 (11), November 2007, 1251-1255.
- [Smith et al. 2005] Smith, B., Ceusters, W., Klagges, B., Köhler, J., Kumar, A., Lomax, J., Mungall, C., Neuhaus, F., Rector A. and Rosse, C. “Relations in Biomedical Ontologies”, *Genome Biology* (2005), 6 (5), R46.
- [Trautwein and Grenon, 2003] M. Trautwein and P. Grenon, "Roles: One Dead Armadillo on Wordnet's Speedway to Ontology", in: P. Sojka, K. Pala, P. Smrz, C. Fellbaum, P. Vossen (Eds.), *Proceedings of the Second International Wordnet Conference (GWC 2004)*, Masaryk University, Brno, Czech Republic, 2003, 341–346.
- [Watson n.d.] Watson Semantic Web Search Engine, <http://watson.kmi.open.ac.uk/WatsonWUI/>. Last accessed July 16, 2008.