Challenges for Ontological Engineering in the Humanities – A Case Study of Philosophy

Pawel Garbacz

Department of Philosophy
The John Paul II Catholic University of Lublin

Abstract. The paper develops an idea of an engineering ontology whose purpose is to represent philosophy as a research discipline in the humanities. I discuss a three recent attempts in this respect with the aim to identify their modelling potential. The upshot of this analysis leads to a new conceptual basis for ontological engineering for philosophy. I show how this framework can be implemented in the form of a simple OWL ontology.

1 Introduction

Representation of knowledge in the humanities poses a number of specific challenges for symbolic Artificial Intelligence. They are mainly due to the idiosyncratic nature of this type of intellectual activity and the features of informational artefacts it provides. Any formal representation thereof needs to account for different, often conflicting, world views adopted by the humanities scholars, the pervasive use of ethnic languages, the instability the technical terminology, in particular the variability of meanings [12]. In addition, the very notion of the humanities, which is to cover all humanistic disciplines, seems to be a grab bag category ([2, p. 222]), which collects rather heterogeneous disciplines like archaeology and performance studies. Thus, despite a number of recent unification initiatives (like TaDiRAH Taxonomy – see: `https://github.com/dhtaxonomy/TaDiRAH` – or Scholarly Domain Model [6]), the prospects of arriving at one common symbolic framework for symbolic knowledge representation are dim.[1]

Some humanistic disciplines seem to have fared better than the others in this respect. For example, there are various formal reference models for library science (Dublin Core, FBFR, etc.) or for cultural heritage (e.g., CIDOC CRM Reference Model – see [5]). Other, like philosophy or musicology, are neglected. In this paper I discuss the challenges for symbolic knowledge representation in philosophy. More specifically speaking, I will investigate the issues and requirements relevant for the ontological engineering paradigm. Section 2 identifies a three recent attempts to representing philosophical knowledge as specimens of the typical solutions to these issues. Section 3 outlines a different type of solution, which aims to overcome some shortcomings of these previous attempts. The implementation of this solution is described in section 4.

[1] All URLs mentioned in this paper were access on May 19, 2015.

Given the state of the art in this domain in what follows I assume that an ontology is a theory rendered in a formal language of logic, usually in a Description Logic language, e.g., in some OWL dialect.

2 Main paradigms in ontological engineering for philosophy

Research in knowledge representation for philosophy is scarce and uncoordinated. This section reports three major attempts at providing an informational artefact to store and reason over the data that come from the philosophical resources. The three specimen in question reveal two important paradigm typologies in ontology development.

One typology concerns the level of detail on which a given ontology represents its domain. One type groups ontologies whose categories grasp only the main differences between entities they refer to. They make as few distinctions as possible and exhibit a relatively small number of categories, which are usually organised in shallow taxonomies. The type at the other end of this typology collects highly discriminative ontologies with numerous categories and distinctions, which usually involve many layers of logical divisions.

The other concerns the amount of the domain knowledge that a given ontology encodes in its structure, which is defined by its terminology (i.e., by its set of terminological axioms in the sense of Description Logic). So there are ontologies designed to include as much of the respective domain knowledge as possible within a given set of expressivity constraints that are imposed on the type of the formal language employed. On the other end of this spectrum there are ontologies that minimise this domain knowledge import.

Although in principle the two dimensions are orthogonal, a discriminative ontology usually involves more domain knowledge than a non-discriminative one. Obviously the above characterisation of these dimensions should be rendered more fine-grained because one ontology may be more discriminative (or knowledge-laden) than another ontology with respect to one group of its categories and less discriminative (or knowledge-free) with respect to another group.

2.1 Indiana Philosophy Ontology project

Indiana Philosophy Ontology (aka: InPhO) project [2] was developed as a support tool for the Stanford Encyclopedia of Philosophy (`http://plato.stanford.edu/`), which is one of the most comprehensive, informative and popular online, open access, and dynamic reference dataset in philosophy.

The InPhO is published as an OWL ontology (of the AL(D) expressivity). The top-most layer is rather sparse – it contains six main categories: Human, Idea, Nationality, Organization, Profession, and Publication. The ontology development focused on the category of ideas, which breaks down into a classification that contains more than 200 subcategories. [2] claims this to be the most noteworthy aspect of this ontology. Its main tenet is "semantic inheritance

relationships holding between the contents of ideas rather than more formal inheritance relationships observed in their types (e.g. social or structural roles)." [2, p. 213]. So, for example, the concept "philosophical idea" is not broken down into kinds of philosophical ideas like concepts, positions, statements, etc., but it is split into philosophical ideas about epistemology, philosophical ideas about logic, etc. Another crucial aspect of InPhO is its approach towards the distinction between (abstract) classes and (concrete) individuals. The InPhO developers chose a pragmatic approach, where a philosophical idea was, as a rule, classified as an individual when it corresponds to an individual entry in the Stanford Encyclopedia of Philosophy.

The InPhO ontology was populated in a semi-automated way, where the initial text mining techniques were audited by the domain experts (i.e., SEP authors) by means of the in-house three-step feedback-harnessing strategy:

1. an expert assesses whether a term found by the standard statistical methods (the tf-idf algorithm, n-gram models, etc.) is relevant for his or her entry;
2. the experts evaluates the level of relatedness of the term for the entry;
3. the expert evaluates the non-taxonomic relationships found in the unvalidated sources (e.g., in Wikipedia).

At the time of writing this paper the Indiana Philosophy Ontology project was still maintained at `https://inpho.cogs.indiana.edu/`.

2.2 PhiloSurfical ontology

The PhiloSurfical ontology [8] was a data component of the PhiloSurfical annotation tool, which was used to contextually navigate a classic work in twentieth century philosophy, L. Wittgenstein's *Tractatus Logico-Philosophicus*.

The developers of the PhiloSurfical ontology focused on the following aspects of the study of philosophy:

1. historical events
2. generic uncertainty
3. information objects
4. interpretation events
5. contradictory information
6. viewpoints
7. varying granularity

As opposed to the InPhO project the PhiloSurfical ontology is built upon CIDOC CRM Reference Model, which is an upper-level ontology used in the cultural heritage systems. The former ontology was extended in a number of directions, including specific types of events related to the philosophical activity. For example, the CIDOC CRM category 'E28 - Conceptual-Object' was extended with a number of subcategories, one of which is the category of philosophical ideas that is defined in terms of the following eight main subcategories: (i) argument-entity, (ii) problem area, (iii) problem, (iv) method, (v) view, (vi)

rhetorical figure, (v) concept, and (vi) distinction. Each of these categories is further specialised into subcategories, e.g., the category of problems is split into 23 subcategories. Therefore, the PhiloSurfical ontology is relatively detailed description of the discipline of philosophy. The OWL ontology available from contains almost 400 OWL classes and more than 300 object properties.

The categories of the PhiloSurfical ontology were used to annotate Wittgenstein's *Tractatus Logico-Philosophicus*, but the results are no longer available at the project's website: `http://philosurfical.open.ac.uk/`.

2.3 Philospace ontology

The Discovery project was aimed at developing personal desktop applications used to enrich the content of Philosource, which is a federation of semantic digital libraries in the field of philosophy [4]. One of the ontologies developed for the sake of this project is the Philospace ontology, which was available from `http://www.dbin.org/brainlets/discovery/ontologies/philospace_0.1.owl`.

The Philospace ontology is a relatively small artefact. Extending another ontology built in this project it contains 21 classes and 20 object properties – see figure 2. All these categories are not specific to philosophy, but appear to be tailored for the needs of representing any humanistic discipline.

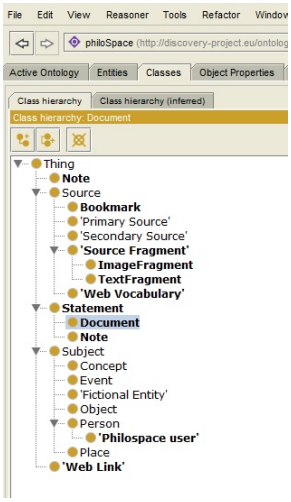

Fig. 1. The class hierarchy in the Philospace ontology

Another ontology developed for the Discovery project, Wittgenstein ontology [9], is currently inaccessible since the two versions of this ontology available at `http://wab.uib.no/wab_philospace.page` contain ill-formed IRIs.

2.4 A faceted typology of engineering ontologies for philosophy

Despite a huge number of classes and their instances the InPhO ontology represents the non-discriminative approach in ontology development. The Idea category contains (as subclasses) philosophical subdisciplines (e.g., logic, ethics, etc.), theories (e.g., bayesianism, connectionism, etc.), arguments (e.g., arguments for the existence of God), and concepts (e.g., mind, causation, etc.). Also within a single category one can find heterogeneous conceptual structures. For example, the category of relations has three instances: medieval theories of relations, relative identity, and logical atomism, where the first is a collection of theories, the second is a concept, and the third is a theory. At the same time it is an example of those ontologies that involve a significant portion of the domain knowledge in its terminology.

	non-discriminative	discriminative
knowledge-free	Philospace	
knowledge-laden	InPhO	PhiloSurfical

Fig. 2. A facet typology of engineering ontologies for philosophy

The PhiloSurfical ontology seems to exemplify the discriminative paradigm in ontology development. The categories it contains are extremely specific. For example, the subclass of (philosophical) problems that deals with problems related to relations is split into problems about the relation of dependence, problems about the relation of independence, problems about the relation of identity, and problems about the relation of difference. Thus, this ontology is an example of ontologies that discriminate between types of conceptual structures they represent and the depth of the distinctions they make depends on the domain they represent. In this case some of the distinctions presuppose the validity of certain specific philosophical views. For example, the class GOD is a subclass of the class SUPERNATURAL-ENTITY – this subsumption is not compatible with those views that see gods as natural objects. Therefore, the PhiloSurfical ontology also involves a significant portion of the domain knowledge.

On the other hand, the Philospace ontology assumes no philosophical knowledge and at the same time belongs to the group of non-discriminative ontologies.

3 Towards a new idea

It seems to me that the discriminative ontologies are likely to be conceptually inflexible and may need to be adapted for new datasets more often. This is due to the high level of detail at which they capture their respective domains. This level of detail may also hinder their reuse as certain conceptual choices made by

their developers may be unacceptable to their potential re-users. At the same time it is this level of detail that results in a more adequate structurisation of the domain knowledge. At the other end of this spectrum we find non-discriminative ontologies. The ontologies of this type are highly flexible and are unlikely to be in the need to adaptation for new datasets. The downsize is now a certain sloppiness in the way they represent their domains. The "one-size fits all" principle seems to be at odds with the aims usually set forth for symbolic knowledge representation. Given these very aims one could expect that an engineering ontology for a given domain should involve, *ceteris paribus*, as much domain knowledge as possible (given the expressivity constraints in question). The reason is that the ontology is, after all, a formal representation of the domain knowledge. Nonetheless, in the case of the humanities, in particular in the case of the discipline of philosophy, this rule needs to be revised. The philosophical research has *not* produced a homogeneous body of consistent knowledge. This field of study is all about controversies, interpretations, and viewpoints. In fact the scope and depth of disagreement is so significant that even the term "philosophical knowledge" looks like an oxymoron. Therefore if we want to maximise the scope of the philosophical research to be represented by an engineering ontology, we need to minimise the impact of particular philosophical assumptions we make building the ontology. Otherwise, we might find ourselves in a position of not being able to express certain views or claims.

For these reasons in what follows I will suggest a conceptual foundation for an engineering ontology for the domain philosophy, which foundation is to (i) maximise the discriminative power of the ontology (ii) while minimising the domain knowledge import. Incidentally, note that any ontology of such kind can fill the empty slot in fig. 2.

The main philosophical assumption is the distinction between texts and text contents, which distinction refines a more familiar distinction between information carriers and information. The notion of *text* is understood here as equivalent to the notion of information content entity from the Information Artefact Ontology (`https://code.google.com/p/information-artifact-ontology/`). So a text is an entity that conveys a certain piece of information. Note that in this sense a text is not a (printed or written) piece of paper, but a certain abstraction of over a collection of such pieces of papers. For example, suppose that each copy of a certain journal paper is a single piece of printed paper. The sense of "text" I use is then that there is just one text (of this journal paper) with multiple copies. The text that these pieces contain is a kind of abstraction over them.[2]

A *text content* is an entity that is existentially dependent on a text in the following sense: when the former exists, then there exists at least one text that expresses the text content. So text contents constitute another layer of abstraction over the pieces of paper we usually refer to as texts. The notion of text

[2] This remark may be rendered in a more rigorous way. Given the principle of abstraction each set of such pieces of papers is an equivalence class of a certain equivalence relation in the set of all pieces of paper.

content is equivalent to the notion of work in the Functional Requirements for Bibliographic Records standard [7].[3]

The role of text contents is to collect texts that convey the same conceptual content, e.g., when one paper contains a text "Endurants exist." and the other paper contains a text "Continuants exist.", then one can arguably claim that these two texts convey the same content.

In order to represent philosophical information I will use three basic types of text contents:

1. categories
2. propositions
3. propositional structures

A *category* is a text content whose role is to represent (i.e., stand for) a collection of entities. Usually categories are rendered in ethnic languages by means of common nouns. The notion of category is understood here rather broadly as it includes also relations (of any arity). So there exists a category of human beings as well as the category of parenthood, the latter representing the particular relations that hold between parents and their children (by means of ordered couples).

A *proposition* is a text content whose role is to represent atomic situations (or, in the philosophical jargon, states of affairs). Examples of situations include that John is a human being, that endurants exist, that no cause follows its effect, etc. A situation is atomic if no part of it is a situation. Usually propositions are rendered in ethnic languages by means of simple sentences. Therefore, propositions are understood as carriers of truth values and objects of the so-called propositional attitudes. So they can be either true or false, accepted or rejected, stipulated of inferred, etc.

Propositions are mereologically complex entities. In particular each proposition contains (as its part) at least one category.

A *propositional structure* is a mereological sum (fusion) of more than one proposition. Usually propositional structures are expressed in ethnic languages by means of complex sentences or sentences concatenations (e.g., as paragraphs, articles' sections and whole articles, books' chapters, etc.). The basic types of propositional structures include arguments and theories.

Note that although propositions and propositional structures are mereologically complex, a category may also contain other categories as parts. For example, the category "vague identity" contains two categories (as its parts): "identity" and "vagueness".

In the humanities text contents are subject to interpretation. The ontology presented in this paper defines the notion of interpretation applicable to categories – interpretation of propositions and propositional structures is a topic of a further study. On the first approximation, an interpretation of a category is a proposition that relates this category to its extension. Different interpretations

[3] I do not find the FRBR distinction between expressions and manifestations useful for the purposes of this paper.

of a category will then define its different extensions. The notion of category extension is systematically ambiguous. For certain (philosophical) categories their extensions are just classes of entities that fall under these categories. Thus, for instance, the category of abstract objects is simply the class of all abstract objects. This group contains those categories whose membership is modally rigid, i.e., it does not depend on temporal indices, possible worlds, contexts, etc. That whether x is an abstract object or not has this characteristics. There is another group of categories whose extensions are modally sensitive. Take any role as an example, e.g., take the category of students. Although John is now a student, he might not be, e.g., he was not and he will not be a student. For such categories the extensions need to be parametrised, either with the temporal or modal aspect (or both). Such extensions may be represented as mappings from these parameters to classes of entities, e.g., the extension of the category of students may be a mapping f_1 from temporal indices to classes of human beings – $f_1(t)$ will be a class of those human beings that are students at time t. Note that in certain research contexts we may need to represent this extension by means of a double parametrisation, e.g., as a mapping f_2, where $f_2(t, w)$ is a class of those human beings that are students at time t and in a possible situation (world) w. The latter may be needed, for example, when one takes into account counterfactual situations in formal models of planning activities. Now the problem with such categories is that although we can enumerate some basic types of extensions, we cannot list all of them as the number and the type of parameters depends on the domain of interest. For example, representing concepts from the psychological point of view we may need take conceptual frames as parameters – cf. [1]. Fortunately, in the case of philosophy most of the categories are modally rigid, so if x is a process, then it is a process in all circumstances.

There is an additional aspect of interpretation for those categories that are relations. Namely, philosophical interpretations of some basic relations as dependence or causality may differ in the arities assigned to these relations. For example, the standard account of the grounding takes it to be a binary relation (e.g., a certain legal fact is grounded in a social fact). But [11] argues for the quaternary account of grounding, where the grounding pattern reads "fact f_1 rather than fact f_2 is grounded in fact f_3 rather than f_4. Even given a fixed arity one interpretation of a relation may differ from the other in the categorical conditions imposed on its arguments. For example, there are two competing accounts of the grounding relation. One imposes no constraints as to what may ground what [10]. The other interprets it as a relation between facts or situations [3]. In order to account for these possibilities my notion of interpretation of relations will take into account (also) the number of arguments for a relation and the categorical restrictions of each argument.

One category may have multiple interpretations which specify the category's different extensions. This implies that we need to construe the former as individual entities, on a par with categories, propositions, and propositional structures. Consequently, the interpretation-sensitivity of categories may be represented by means of the *ternary* relationship that binds category interpretations, interpreted

categories, and categories' extensions. In the case of relations it may be handy to be able to employ the quaternary relationship that binds relation interpretations, interpreted relations, relation arguments' indices, and their categories.

4 Implementation

The above ideas were materialised in the form of an OWL ontology, by the name of OntOfOnt. The ontology is scoped to only one philosophical discipline: metaphysics (ontology). In other words, it is designed as an engineering ontology for philosophical ontology.

As of the time of writing the ontology has 39 classes, which are related by 39 object properties. The ontology's axiomatisation attempts to employ the full expressive power available in OWL 2 DL, so its expressivity is ALCROIQ(D). Still, a more adequate formal characterisation of the notions discussed in the previous section would require the full strength of the first order logic. The class hierarchy is shown in fig. 3.

Since OntOfOnt is not a foundational ontology, its upper level categories were "stitched" up to the categories of two foundational ontologies: CIDOC CRM and Information Artefact Ontology - see table 1. The various versions of the OntOfOnt ontology are available from the URLs specified in table 2.

OntOfOnt	CIDOC CRM	IAO
agent	owl:equivalentClass E39_Actor	owl:subClassOf BFO_0000030
category_extension	owl:subClassOf E70_Thing	owl:subClassOf BFO_0000141
membership	owl:subClassOf E70_Thing	owl:subClassOf BFO_0000141
intentional_entity	owl:equivalentClass E73_Information_Object	owl:subClassOf BFO_0000031
intentional_entity_carrier	owl:subClassOf E24_Physical_Man-Made_Thing	owl:equivalentClass IAO_0000030

Table 1. OntOfOnt upper level categories characterisation

OntOfOnt version	URL
base ontology	`http://www.metaontology.pl/metaontology.owl`
base ontology embedded in IAO ontology	`http://www.metaontology.pl/iao_metaontology.owl`
base ontology embedded in CIDOC-CRM ontology	`http://www.metaontology.pl/cidoc_metaontology.owl`
base ontology populated with `http://philpapers.org` data	`http://www.metaontology.pl/metaontology_populated.owl`

Table 2. OntOfOnt versions

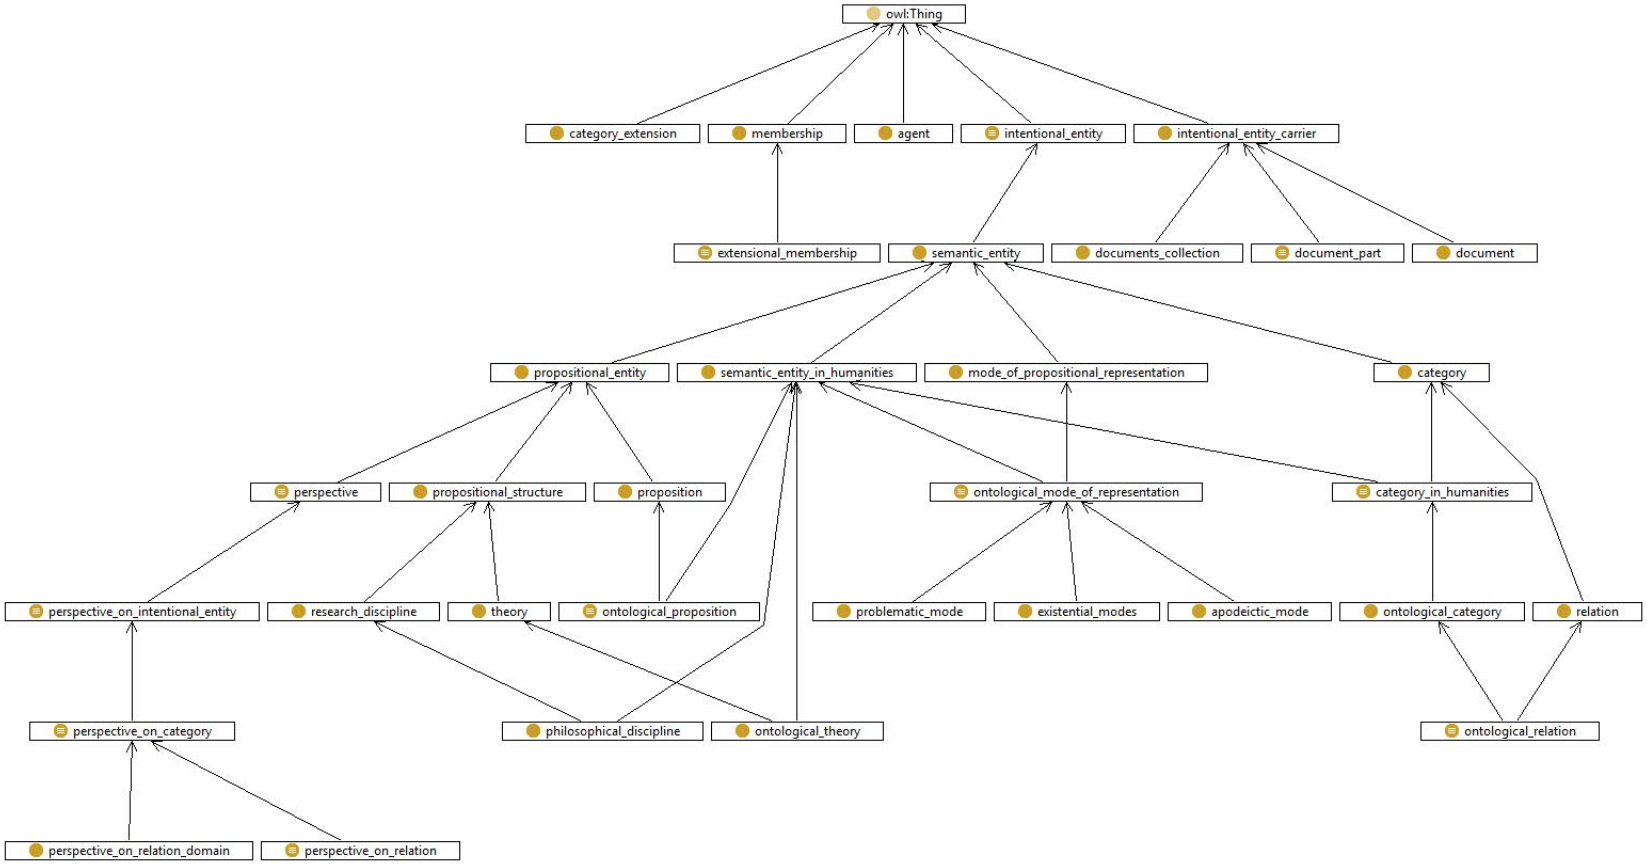

Fig. 3. OntOfOnt class hierarchy

The relationships that support the representation of category interpretations cannot be directly expressed in OWL languages, which are restricted to binary relations. In order to overcome this restriction I employed the standard procedure of reification, whose results are shown in fig. 4.

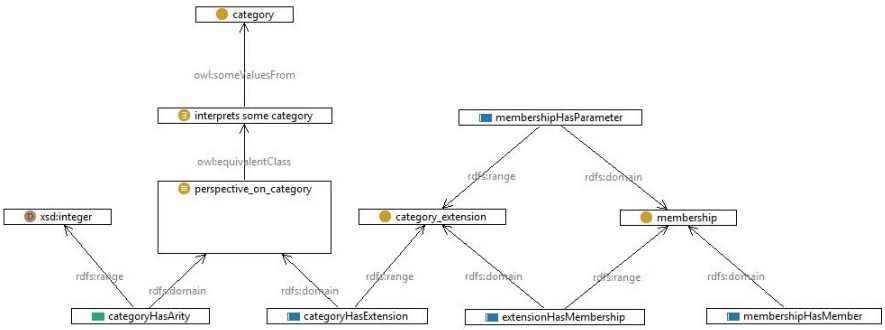

Fig. 4. OntOfOnt reifications

In order to illustrate its modelling potential the ontology was populated with the data available from the `http://philpapers.org` website. This website organises the scholarly papers in philosophy with respect to a number of categories. In particular it defines a set of categories relevant for philosophical ontology. A close look at `http://philpapers.org`'s categories revealed that they include all three types of text contents defined above. These categories and the papers assigned thereof were loaded to OntOfOnt accordingly. In addition during the load a number of additional categories were identified on the basis of the papers' abstracts. The identification process was semi-automated by the user's selection of the keywords with the highest tf-idf ranks. The whole load process was supported by a purpose-specific JAVA application. Note that this process did not involve any philosophical analysis, so it assumed, among other things, that each category has the same interpretation throughout the papers assigned to it.

5 Further Work

The main theoretical issue to be resolved has to do with the impoverished notion of category interpretation assumed in this paper. The notion described above grasps only the extensional aspect of categories, ignoring their intentional dimension. In a sense all categories with the same extensions are equivalent. So, for example, the category of equilateral triangles will be characterised exactly in the same way as the category of equiangular triangles despite the fact that these

categories are different because they are defined in different ways. A more comprehensive notion of category interpretation needs to take this essential aspect of categories into account.

As far as the applicative issues are concerned the main obstacle again hinges upon the notion of interpretation. This time the task is to specify for each text content the interpretations involved in this content and identify the same interpretations across different contents. To this end we need laborious philosophical analysis of the respective resources. As a rule, this process cannot be automated or even computer-aided, but it needs to be done by a human user who is sufficiently conversant in the domain at stake so that he or she could identify and distinguish between different interpretations of thereof. Only then the full expressive potential of OntOfOnt can be employed and the specific nature of philosophy as a discipline in the humanities can be established.

References

1. Barsalou, L.W.: Frames, Concepts, and Conceptual Fields. In: Lehrer, A., Kittay, E.F. (eds.) Frames, Fields, and Contrasts, pp. 21–74. Lawrence Erlbaum Associates (1992)
2. Buckner, C., Niepert, M., Allen, C.: From encyclopedia to ontology: Toward dynamic representation of the discipline of philosophy. Synthese 182(2), 205–233 (2011)
3. Correia, F., Schnieder, B.: Grounding: An opinionated introduction. In: Metaphysical Grounding, pp. 1–36. Cambridge University Press (2012)
4. DIorio, P.: Discovery. D1.8 Final Report. Tech. rep. (2009), `http://www.discovery-project.eu/reports/discovery-final-report.pdf`
5. Doerr, M.: The CIDOC Conceptual Reference Module: An Ontological Approach to Semantic Interoperability of Metadata. AI Magazine 24(3), 75–92 (2003)
6. Hennicke, S., Gradmann, S., Dill, K., Tschumpel, G., Thoden, K., Morbindoni, C., Pichler, A.: D3.4 Research Report on DH Scholarly Primitives. Tech. rep. (2015), `http://pro-beta.europeana.eu/files/Europeana_Professional/Projects/Project_list/Digitising%20Manuscripts%20to%20Europeana/Deliverables/D3.4%20Research_Report_on_DH_Scholarly_Primitives.pdf`
7. IFLA: Functional Requirements for Bibliographic Records: Final Report. K. G. Saur (1998)
8. Pasin, M., Motta, E.: Ontological requirements for annotation and navigation of philosophical resources. Synthese 182(2), 235–267 (2011)
9. Pichler, A., Zöllner-Weber, A.: Sharing and debating Wittgenstein by using an ontology. Literary and Linguistic Computing 28(4), 700–707 (2013)
10. Schaffer, J.: Metametaphysics: new essays on the foundations of ontology, chap. On What Grounds What, pp. 347–383. Oxford University Press (2009)
11. Schaffer, J.: Grounding, transitivity, and contrastivity. Metaphysical grounding: Understanding the structure of reality pp. 122–38 (2012)
12. Veltman, K.H.: Towards a Semantic Web for Culture. Journal of Digital Information 4(4) (2006)