CHAPTER 11

PhiloSurfical: an Ontological Approach to Support Philosophy Learning

Michele PASIN[a,1] and Enrico MOTTA[a]
[a] *Knowledge Media Institute, The Open University, UK*

Abstract. As the Semantic Web is increasingly becoming a reality, the availability of large quantities of structured data brings forward new challenges. In fact, when the content of resources is indexed, not just their status as a text document, an image or a video, it becomes important to have solid semantic models which avoid as much as possible the generation of ambiguities with relation to the resources' meaning. Within an educational context, we believe that only thanks to these models it is possible to organize and present resources in a dynamic and contextual manner. This can be achieved through a process of narrative pathway generation, that is, the active linking of resources into a learning path that contextualizes them with respect to one another. We are experimenting this approach in the PhiloSurfical tool, aimed at supporting philosophy students in understanding a text, by presenting them 'maps' of relevant learning resources. An ontology describing the multiple aspects of the philosophical world plays a central role in this system. In this chapter we want to discuss some lessons-learned during the modeling process, which have been crystallized into a series of reusable patterns. We present three of these patterns, showing how they can support different context-based reasoning tasks and allow a formal conceptualization of ambiguities that are primarily philosophy-related but can be easily found in other domains too. In particular, we describe a practical use of the ontology in the context of a classic work in twentieth century philosophy, Wittgenstein's *Tractatus Logico-Philosophicus*.

Keywords. Philosophy, ontology, digital narratives, semantic web, Wittgenstein

Introduction

The need to specify and separate the information about the context of usage of a learning resource, from the resource itself, is one of the main reasons behind the creation of various kinds of metadata schemas. In the past years, this work has focused around the notion of learning object (LO) [1], as the technology capable of guaranteeing interoperability to the rapidly growing number of Web-based educational applications. However, increasingly researchers are now arguing that LOs' metadata are not fine-grained enough to non-trivial composition of resources, e.g. when constructing a curriculum [2]. As a result, as attested by a series of workshops held worldwide [3], the e-Learning research community has begun looking at the potential

[1] Corresponding Author.

for e-Learning of the emerging Semantic Web technologies. In this context, ontologies [4] have been proposed by many [5,6] as a technology that can be used to *complement* the functionalities of traditional LOs metadata standards – for example, ontologies can be useful when it is important to describe precisely and unambiguously a specific domain of interest, for this is the main *content* a LO is dealing with; or to the end of producing an extensive formalization of learning and teaching strategies, so to have much more control over the *context* the learning resources are being used within [7].

Our approach, in compliance with this new research direction, supports the enhancements of LOs' metadata through the usage of ontologies so to represent the content of learning resources at a finer level of detail. More precisely, in the PhiloSurfical tool this approach is realized through the formalization of a humanistic discipline, *philosophy*. An ontology to describe and organize *theories*, *schools of thought*, *arguments*, *problems* and their relations to other philosophical concepts will allow the annotation of the learning material, and, subsequently, its dynamic reorganization with a degree of accuracy and flexibility well beyond the one provided by standard LOs metadata.

By doing so we aim at providing a platform that supports philosophy students in understanding key aspects of the discipline's discourse. This is achieved by means of a *pathway* creation process, i.e., an approach that gives students the means for *contextualizing* the resources they have found so to better analyze and interpret them in the light of the multiple roles they play in the world of philosophy. Our approach takes the notion of a *learning pathway* as a "system of specially stored and organized narrative elements which the computer retrieves and assembles according to some expressed form of narration" [8] and attempts to transpose it within the specific scenario made up of philosophical entities.

For example, we can think of a young philosopher trying to understand Wittgenstein's *picture theory of language*. Our purpose is to let her discover the significance of this theory by putting it into different perspectives (i.e., the pathways) and autonomously exploring how it relates to other philosophical entities (e.g., theories, events or people). Among the pathways available our student could find the following ones:

- the *critical explanation* of a theory (a meta-historical learning path that highlights the opposing theories, and the problems on which they are focused),
- the *historical contextualization* of a theory (a learning path that shows associated information about an author, or the historical period, or other contemporary important theories in different research areas),
- the description of the *whole body of work of an author* (a learning path that recollects all the activities and results of an author, and organizes them according to the user's preferences),
- the *intellectual lineage* of a concept/theory (a learning path that follows the influence of ideas throughout the history of thought, across different areas and historical periods).

The aim of this chapter is twofold. Firstly we consider three ontological lessons-learned which emerged as fundamental in supporting philosophical learning pathways. Under this respect, the modeling patterns we are presenting are quite different from the patterns discussed in other works such as that one of Gangemi and colleagues [9],

where the focus is more on the architectural issues involved in the ontology creation process. In particular, the patterns we are describing in the next sections represent some modeling decisions that are meant to *guide the interpretation* of philosophical texts, so to have formal models that are *applicable* for providing non-trivial navigation mechanisms. We believe that such a modeling can have a significance that goes beyond the specific domain of philosophy and can be reusable within more generic areas of interest.

Secondly, we show the reader a practical application of this ontology, the PhiloSurfical tool. This is a web application that supports users in learning about a philosophical text, Wittgenstein's Tractatus Logico-Philosophicus [10]. By relying on the multiple representations of our philosophical ontology, PhiloSurfical's learning pathways lets students benefit from multiple perspectives on the text and on related resources. For example, they can reorganize the text according to the relevance of a single annotation, e.g. the concept of "logical-independence", or they can adopt more complex strategies to retrieve other resources that are not directly related to the Tractatus. In particular, in the following sections we are going to discuss the functioning of a specific type of pathways, the *theoretical* ones.

The rest of the chapter is organized as follows: section 1 provides some pedagogical background about philosophy learning; section 2 gives an overview of the ontology we built for representing the world of philosophy, for then discussing the details of three important modeling patterns that emerged during the work; section 3 describes PhiloSurfical, our narratology-inspired prototype application. Finally, section 5 concludes the chapter by providing information about other related work.

1. Constructive Learning in Philosophy

Let us think again about our philosophy student, while she is tackling Wittgenstein's theory of language, as exposed in the Tractatus Logico-Philosophicus. She would probably read Wittgenstein's texts several times, analyze the language being used and deepen her understanding of a number of concepts the philosopher's argument relies upon. Also, she would likely make use of other reference material about this topic, so to gain insight into the historical and theoretical contexts the theory originated from. Is it addressing a long-standing problem, or does it raise a completely new one? Who has been influencing Wittgenstein, and how much of his ideas can be related to other preexisting philosophical work? These are the type of questions we expect our student to try to answer.

At a more general level, we could say that our student is *actively exploring* this new philosophical 'territory'. An active style of learning, not just, for example, a passive reading and remembering of what is read, is reputed by many as being the main cause for successful learning. In educational theory, this thesis (and others related to it) is one of the central tenets of doctrines such as constructivism [11] and situated cognition [12]. Their importance and academic relevance, beyond the various and inevitable debates, is widely acknowledged.

For example, an active style of learning implies that, when facing a text, although a teacher's explanation is of help in the learning process, he/she is not the main reason for it. In fact, according to this position teachers are more often viewed as 'knowledge facilitators', in opposition to the traditional figure of the 'knowledge dispenser'. In general, students are advised to engage directly with a subject matter (e.g., an author's

text), in order to obtain their own understanding and actively *construct* a meaning out of it.

However, this picture is quite a simplified one. While an active style of learning is relatively easy to foster in natural, everyday situations (for example, when learning how to ride a bike or how to speak a language), this is not the case for the more artificial, academic learning. The learning and teaching of philosophy can be taken as an example of this difficulty. Philosophy, as other subjects such as theoretical physics, mathematics and logic, deals only with abstractions. That is, in Laurillard terms, "descriptions of the world" [13]. As a consequence it is harder to situate its learning in a natural context and it is also hard to apply constructivist approaches to teaching.

In such an academic and abstract context, what are the ideal students' activities which can lead to a successful learning experience, and what are the best methods and situations to support them, is the object of much debate [14-16]. But even if a general agreement on this matter will hardly be reached, we can still attempt to define some essential requirements to achieve in the context of philosophy teaching. More precisely, we agree with Carusi that the three most important skills to develop in a philosophy student must be (a) analysis, (b) argument and (c) interpretation. As the author remarks, the "three skills are interwoven as analysis requires interpretation, and argument depends on the prior abilities to analyze and interpret correctly other philosophical positions" [17].

In particular, in table 1 we detail Carusi's lengthier description of what each of the skills may entail, as far as the student is concerned.

Table 1. The three major philosophical skills (from Carusi, 2003)

Skill	Description
Analysis	• analyze a philosophical problem or position into its component parts and be able to tell how they are connected together; • analyze an argument into premises and conclusions, and reconstruct the structure of the argument, filling in implicit premises where necessary; • analyze philosophical texts into sections and be able to see the connections between sections.
Argument	• understanding of the standard fallacies; • being able to distinguish between inductive and deductive arguments, and being able to say what constitutes an acceptable argument of both kinds; • understand the role of counter-examples and be able to use them; • understand the role of analogies and be able to use them; • understand the role of thought experiments and be able to use them.
Interpretation	• Interpretations should be coherent in that they should not contain inconsistencies or contradictions. • Interpretations should be cogent in that they should account for as much of the text as possible within a unified framework. • Interpretations should be informed by an understanding of the historical tradition in which the text is embedded and the meanings of concepts and terms as specified within that tradition. As a minimum, this should include some knowledge of history of ideas in philosophy.

With the PhiloSURFical tool (see section 3) we aim especially at supporting the (a) analysis and (c) interpretation skills development, through an environment which allows constructing advanced strategies to present annotated resources to the user, in the form of browsing facilities and narrative generation. The active involvement of the student in a process of semi-structured navigation (the structure being provided by the ontology) guarantees her engagement with the subject matter in a constructivist manner.

2. Engineering Philosophical Knowledge

2.1. Overview: an event-centered design

The specific approach used to realize the PhiloSurfical ontology has at its centre the decision to employ the CIDOC Conceptual Reference Model [18] as a starting point for our formalizations. The CRM ontology is a renown ISO-standard which aims at supporting semantic interoperability for museum data. In the following sections, we are referring to version 4.2 of the ontology [19].

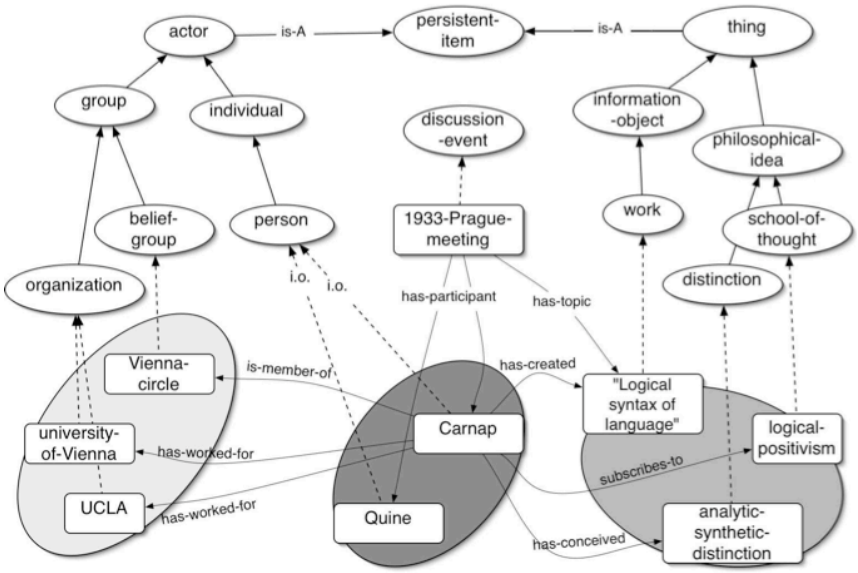

Figure 1. Example of an event-based representation

The choice of using the CRM was motivated by two reasons. Firstly, for its widely recognized status as a standard for interpreting cultural heritage data. In fact, by reusing and extending an existing and internationally recognized ontology, we will give our tool's users more chances to benefit from the emerging Semantic Web infrastructure. Secondly, for its extensive event-centered design. This design rationale, in fact, appeared to be appropriate also when trying to organize the history of philosophy: even if it is common to see it as an history of ideas, stressing the importance of the theoretical (i.e. meta-historical) dimension, this cannot be examined without an

adequate consideration of the historical dimension. That is, a history of the events related (directly or indirectly) to these ideas.

As an example, in figure 1 we can see an event-centered representation in the PhiloSurfical ontology. The `persistent-item` class, which is one of the five classes composing CIDOC's top layer (together with `time-specification`, `dimension`, `place` and `temporal-entity`) subsumes `thing` and `actor`. The two branches of the ontology departing from them can have various instances, which are related by taking part (in various ways) to the same event ("1933-Prague-meeting"). This kind of modeling, in the context of the PhiloSurfical tool, is extremely useful because of the multiple navigational pathways it can support (e.g. we could move to another event having the same topic, or to another topic treated during the same event, etc.). Please note that in the figure some relations (e.g., *has-worked-for*) are graphical shortcuts for the actual and lengthier formalization of the relevant event (e.g., an `event` instance stating that an actor worked for an institution at some point in time etc.).

From the implementation point of view, the ontology has been prototyped by using the Operational Conceptual Modelling Language (OCML) [20], which provides rich support for both specification and reference. Import/export mechanisms from OCML to other languages, such as OWL and Ontolingua, ensure symbol-level interoperability. Please notice that in the next sections we used different fonts depending on whether we refer to classes in the ontology (e.g., `event`) or properties associated to them (e.g., *has-duration*). Instances are always double quoted (e.g., "the concept of will"). For that regards the figures, classes are oval-shaped, rounded rectangles stand for instances and arrows represent relations. In particular, if not labeled otherwise, dashed arrows stand for the *instance-of* relation, while solid arrows stand for the *subclass-of* relation.

At the time of writing, the ontology[2] counts 348 classes, partly integrated from other relevant semantic models and partly identified through various knowledge acquisition techniques (formal and informal). In conclusion, it is worth remembering that our ontology resulted as being the first and most ambitious attempt to provide a formal meta-language usable for describing the world of philosophy. Although we used it mainly in the context of Wittgenstein's philosophy (see section 3), the ontology is very abstract and could be easily applied to other philosophical domains too. We provide an extensive description of all of its features in another publication [21].

In the next sections we will present three ontological issues we encountered during the modeling process, together with the solutions we contrived in order to solve them. As in the example above, the derived modeling patterns aim at taking advantage of the multiple meanings a philosophical entity (e.g. an idea, a text or an event) can have, by making these meanings explicit and employable when building novel exploration mechanisms. In other words, according to our approach 'ambiguities are good' because, if properly identified, they let us explore the domain in different and interesting ways.

2.2. Pattern #1: is rationalism a school of thought or an event?

The first pattern originates from the fact that in our everyday language we refer to belief groups, intellectual movements and schools of thought ambiguously, often using the same word. For example, let us consider the following three statements:

[2] The ontology is available online at http://philosurfical.open.ac.uk/onto.html

a) “Throughout history, the attacks of *rationalism* against empiricism has diminished”
b) “Descartes was one of the founders of modern *rationalism*”
c) “This theory is clearly a new and re-shaped *rationalism*”

Initially, we set out to model concepts such as “rationalism” by adding a philosophy-specific subclass to CIDOC’s `period`. In fact, according to CIDOC, `period` (which is a direct subclass of `temporal-entity`) should subsume prehistoric or historic periods, or even artistic styles. This is motivated by the fact that "it is the social or physical coherence of these phenomena that identify a Period and not the associated spatio-temporal bounds" [19]. This seemed to apply quite neatly also to cultural and philosophical periods, thus we have added `intellectual-movement` and its subclass `philosophical-movement` to the hierarchy.

However, at a deeper ontological analysis, we came to the conclusion that in the sentences above we are using the same word to express three different meanings. Precisely, in a) “rationalism” is the label referencing to a *group of people*, in b) we are meaning an *event*, while in c) we are probably referring to an *abstract idea*.

A modeling pattern (figure 2) achieves the goal of expressing both the difference in meaning and the interrelations of the three senses implied by words such as “rationalism”. This pattern involves subclasses of `actor`, `period` and `view` (a type of abstract philosophical idea, as we shall see later, expressing a standpoint). The ambiguity of a term such as "rationalism" can be clarified, since the semantic model keeps the three different ways to intend the word into a consistent representation. By doing so, we are providing a context of usage for such ambiguous terms, and a direct way to navigate coherently among entities that are ontologically quite distinct (i.e. from `temporal-entity` to `actor` and `propositional-content`, which belong to separate branches of the ontology). Moreover, such a context-specification could be used for by a reasoner to derive inferences from incomplete or inconsistent data sources, or for performing information extraction.

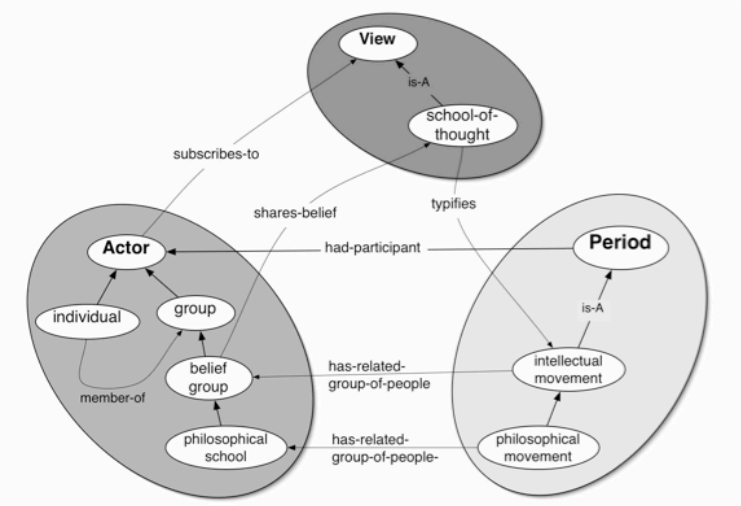

Figure 2. The actor-event-view modeling pattern

So, for example, we can describe the “enlightenment movement” in the following way[3] (note that the temporal relations are specified here as slots, but are usually inferred whenever the appropriate time specifications of the other periods were provided):

```
INDIVIDUAL Enlightenement
    instance of:                        Intellectual-movement.
    has-time-specification :            18th-century.
    overlaps-in-time-with :             scientific-revolution,
                                        renaissance.
    meets-in-time-with :                french-revolution,
                                        american-revolution,
                                        romanticism.
    overlaps-with :                     age-of-Reason,
                                        neo-classical-art.
    took-place-at :                     germany,
                                        france,
                                        britain,
                                        spain.
    has-related-group-of-people :       enlightenment-group-of-people.
    is-typified-by :                    enlightenment-conception.
```

The last two slots in the formalization above have a special importance, for they serve the purpose of interrelating the three different senses highlighted in the pattern. In particular, the slots *has-related-group-of-people* and *is-typified-by* link the “enlightenment” instance (an `intellectual-movement`) to the relevant instances of `group-of-people` and of `school-of-thought`.

2.3. Pattern #2: not all views are theories!

The second pattern is related to the fact that people often employ the term ‘theory’ in a loose manner, over-classifying views with different characteristics. Consequently, a thorough formalization of these entities proved to be an important meta-model usable by students for ‘learning the differences’ among the different theory-types and relationships.

In our ontology, `view` has been defined as a generic class referring to philosophical ideas expressing a viewpoint. That is, propositions picturing a perspective on the world in the form of more or less structured interpretations of things and events. Examples of `view` are "solipsism", "theory of evolution by natural selection", "philosophy of Plato" or "a name has a meaning only in the context of a proposition" (i.e. Frege's context principle). Because of their ‘categorical’ attitude, views usually define concepts and, in general, create the context for the definition of other meanings too (e.g. problem-areas, problems, methods etc.). A number of properties connect views to the other philosophical-ideas: e.g. views can use other ideas, tackle problems, influence and support/contrast each other, be-supported by arguments. However, the feature we want to highlight here is how views can have different granularities: from our analysis of the literature, we identified four of them. This classification is mainly related to the degree of generality they exhibit, and the level of complexity they have. So, we can have (as shown in figure 3):

[3] Although OCML has a simple frame-like syntax, in order to facilitate readability here we are using an abstracted syntax.

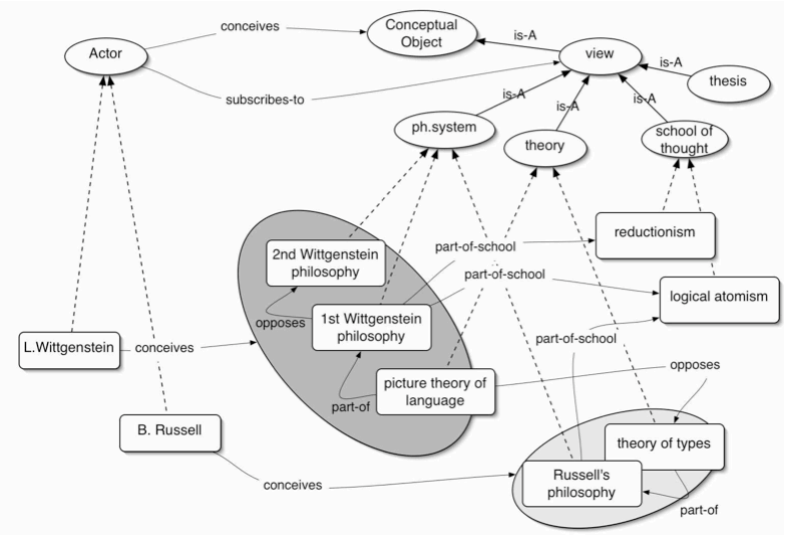

Figure 3. The view-types instantiation

- `Thesis`: it is the least structured `view`, as sometimes it consists only of a standpoint in the form of a statement (i.e. an assertion). So, for example, in the context of Wittgenstein's "picture theory of language", a thesis can be the "independence of the state of things".

- `Theory`: it is a systemic conceptual construction with a coherent and organic architecture. A `theory` explains a specific phenomenon (or a class of phenomena) and typically answers to an already existing problem. Examples can be Darwin's "theory of evolution" or Quine's "verification theory

- `Philosophical-system`: it might appear as a `theory`, at first sight, but it differs from it essentially for its generality. That is, because it spans over various problem-area, while a theory is usually confined to one problem-area only. As a consequence, `theories` are usually *part-of* `philosophical systems`. We can therefore define a system as the set of a person's views that are consistently connected to each other, in such a way to form a unity (in a way, this class refers to what is normally called the "philosophy" of a thinker).

- `School-of-thought`: this class refers to the set of theory-types, or generic standpoints, which in the history of thought have acquired a particular significance and, seemingly, a life on their own. They correspond to widely known conceptions, or standardized intellectual trends that hint at typical ways to answer a problem (or a set of problems). Examples are "pacifism", "animism", "expansionism", "empiricism" or "monism". A `school-of-thought`, compared to the other `views`, is not as formalized and specific as a `theory`, and not as general and systematic as a `philosophical-system`.

Thanks to this quadruple classification, it is possible to specify all the hierarchical and mereological relationships among views with a good degree of precision. From the point of view of learners, this modeling pattern facilitates the creation of pathways that place a theory or school of thought within the larger *theoretical* context, i.e., showing how it is related to to other intellectual entities.

2.4. Pattern #3: 'problematic' problem areas

The third pattern we are presenting wants to provide a way of expressing the distinctive features of 'fields of study'. It is normal for philosophy learners to refer to the topics they are studying in terms of the *area* they belong to, e.g., *metaphysics*, *logic*, *philosophy of language* etc. Indeed, the way the discipline is organized reflects some common denominators of philosophical research: these can be some classic problems investigated in philosophy, or the generic approaches used to solve them. In order to provide explorative pathways that focus on these particular aspects of philosophical discourse we created a modeling pattern centered around these notions.

As we will see, one the major difficulties here arises from the fact that we can interpret fields of studies in at least two different ways: a generic one (e.g. the field of "physics") and a specific one (e.g. "Newtonian physics"). The pattern models the relations between them.

Our starting point is a *problem-centered* approach, that is, the decision to see the activity of philosophers as essentially an ongoing process of specifying and giving solutions to problems. Consequently, we consider any recognized area of study, of whatever type or dimensions, as a `problem-area`. In its simplest version, a problem-area is composed by a set of problems linked by different relational schemas, but in general, tying around a main theme. This theme, in our ontology, can be represented through a `problem` (*has-central-problem* property) or thanks to a `thesis` functioning as a criterion (*has-criteria property*). For example, "psychology", when treated as a problem-area, can gather problems tied to the "mind-definition" problem, to the problem of "relating human behavior to brain activities", or to the thesis that "brain and mind can be investigated with the methods of natural sciences".

Other features of problem-areas are that they can be related to each other (e.g. "mathematics" and "philosophy of mathematics") and that they can be organized into simple hierarchies (e.g. "internet-ethics" is a sub-area of "ethics"). However, we realized soon that "psychology" has a role and significance in our world that goes beyond a mere problem area. In a similar fashion, "ethics" or "cognitive science" would not be properly characterized only as instances of problem-area, for they also refer to theories or methods which have become intrinsically related to the definition of the area.

Moreover, if we consider the history of thought, the topic and description of problem areas have always been subject of many debates: different views aspire at having the ultimate vision about what the central issues to look at are, or the right methods to take. In this respect, problem-areas are not very different from other ideas that can be defined by multiple views. For example, we can just consider how different was the sense given to "philosophy of language" by the first philosophy of Wittgenstein and the second one.

In order to catch these subtle differences, we defined the class `field-of-study` as a `problem-area` that has been socially and historically recognized as separate from the others (and from being a mere agglomerate of problems). In the ontology, this is reflected by the fact that a `field-of-study` is not just specified by a criteria, but is *defined-by* a `view`. It is also characterized by the fact that it collects not only problems, but also ways to solve or tackle them (i.e. theories and methods). The distinguishing properties are therefore *defined-by-view*, *has-exemplar-theory* and *has-methodology*.

Finally, a last tricky issue regarding fields of study must be addressed. This does not emerge when treating relatively isolated entities such as "phrenology", but it clearly

is an issue if we consider, say, “physics”. In our everyday language, and also in the organization of academic programs, we usually refer to “physics”, “psychology” or “philosophy of mind” as *generic* fields of study. What this means, is not really clear. In fact, when we delve into them (or even more, if we ask for clarifications to a practitioner), we discover quickly that there are many “physics”, “psychologies” and “philosophies”, at least as many as the views defining them. From our ontological perspective, these would all be separate instance-candidates of the `field-of-study` class. However, we also need to represent the fact that they are all part of a more generic (and probably emptier, for that regards its meaning) type of field of study.

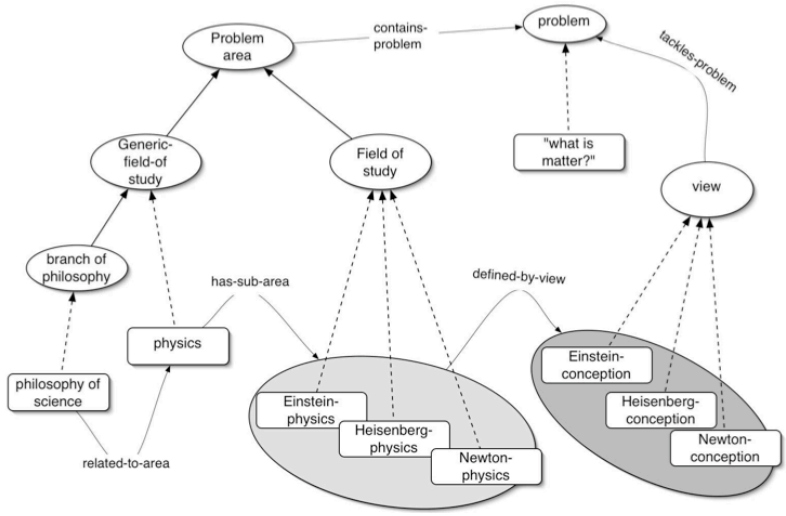

Figure 4. Problem areas and fields of study

Our solution to this problem consists in the creation of a `generic-field-of-study` class, which has no defining view but the views defining the specific fields-of-study that are claimed to be part of it. In other words, we are formalizing the fact that generic fields of study such as “physics” or “philosophy” can be defined only *extensionally*. So:

```
CLASS Generic-Field-of-Study
     subclass-of :Problem-area.
     defined-by-view :
                    :range View.
                    :range-constraint:          [?GF defined-by-view ?V]
                                          =>
                                        Exists ?F:Field-of-Study
                                                [?GF has-sub-area ?F]
                                                [?F defined-by-view ?V].
```

In the formula, the variables ?GF, ?V and ?F refer respectively to `generic-field-of-study`, `view` and `field-of-study`. Therefore, doing so we can maintain the interoperability between specific thinkers’ definitions of classic problem areas, and the generic but useful ways to refer to them. In figure 4 we give a graphical overview of this modeling pattern, highlighting the important relationships among the classes involved. Please note that in this figure we used a graphical ‘shortcut’: when a relation is attached to a group of instances, that is to mean that the relation is repeated over all of those instances. For example, the `generic-field-of-study` instance “physics” exhibits

the property *has-sub-area* three times, corresponding to the three instances of `field-of-study` we grouped together.

3. The PhiloSurfical Tool

3.1. Overview

In order to test the usage of the ontology within a specific philosophical scenario we created PhiloSurfical[4] (see fig. 5). This is an application that supports learning about a classic work in twentieth century philosophy, Wittgestein's "Tractatus Logico-Philosophicus" [10].

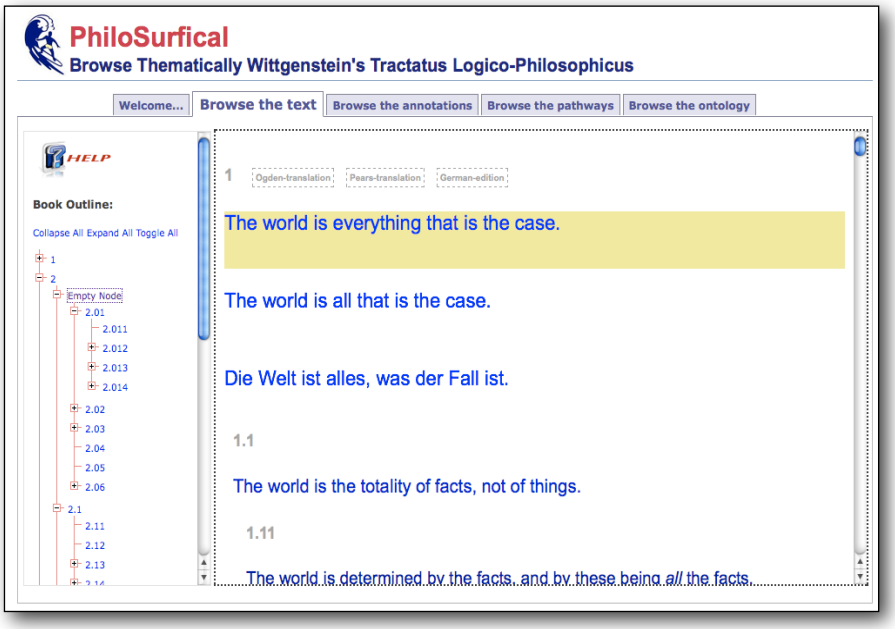

Figure 5. Screenshot of the PhiloSurfical application

The PhiloSurfical's tool functionalities, and in general, the envisaged context of usage which has been guiding the ontology engineering process is the following: the semantic model should support the *reconstruction* of the history of ideas, by relying on structured information about the *practical* domain and the *theoretical* domain of thinkers. For example, within an educational scenario where young philosophers try to understand domain notions (in a wide sense, comprising ideas and events), these functionalities will exist in the form of mechanisms for contextual navigation and linking of relevant resources. As a result, we expect such a service to facilitate the discovery of (related) unknown resources, which can be used by students and scholars during the process of answering difficult problems.

This methodology, which has been previously defined as *ontology-based* navigation [22], can be further developed by means of an approach modeled on

[4] The application is available online at http://philosurfical.open.ac.uk/

narratology [23]. As already discussed in an earlier publication [24], following structuralist theorists we can sketch out the structure of a *narrative* as the union of a *story* (what is told) and a *discourse* (the 'how' of what is told, that is, the specific way in which the basic elements of a story are re-organized and conveyed to the listener, in order to create different effects).

In our narratology-inspired approach, a formal ontology can be used to express the semantics of the different elements composing a story, so that it is also possible to formalize the way a discourse recomposes the same elements according to different criteria. So, for example, the same chosen set of 'atomic' philosophical events could be ordered following a *historical* perspective, a *geographical* one or even a *theoretical* one. Similarly, the same set of philosophical ideas could be organized differently if investigated under a *problem-centered* perspective, a *theory-centered* one, or simply one based on their *historical succession*.

3.2. Knowledge base creation

It is important to remember that although one of the aims of the ontology was to facilitate data-exchange among distributed resource-providers, for bootstrapping purposes (as the availability of free and adequately encoded philosophical data on the web is still limited), PhiloSurfical strongly relies on an internal knowledge base of our creation. As suggested by recent projects such as the DBPedia [25], or the Discovery project [26], we envisage that in the near future this situation will change, as much more structured data about philosophy will be made available.

Our knowledge base was constructed in three phases: first, we transformed the text itself in a format compatible with our ontology - i.e., we instantiated classes representing the text *and its paragraphs*. Second, we annotated the text by working in collaboration with a Wittgenstein scholar, Andrea Bernardi; in this phase we instantiated classes representing *ideas and relations among ideas*, and *indexed* the text using these representations. Third, we enlarged the knowledge base by 'scraping' philosophy-related information from various websites in the public domain; in particular, we created more instances of *philosophers and philosophical schools of thought*.

At the end of of this process, we gathered a total of more than 20 thousand instances connected to Wittgenstein and his philosophy. It was not our purpose to create an exhaustive resource about the Tractatus; accordingly, we stopped refining the knowledge base as soon as we thought we had reached a critical mass of data, usable for testing the ontology through our 'learning pathways' approach.

3.3. System description

From the technical point of view, PhiloSurfical is a *lisp* web-application running on the Lispworks environment [27]. It uses OCML [20] for the knowledge representation and storage functionalities, and Hunchentoot [28] as a web-server.

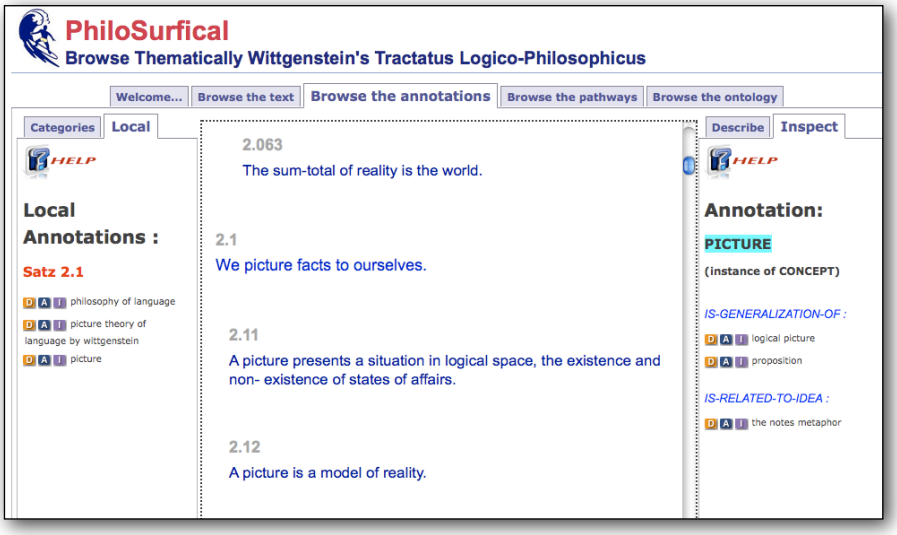

Fig. 6. The 'Browse the annotations' tab in PhiloSurfical

The application is organized into five sections or tabs. We attempted to organize the tabs' sequence according to their increasing difficulty of usage (namely, the first tab requires less learning effort than the second one, the second one less than the third one, etc.). By doing so, we wanted users to have a more gradual encounter with the software. This becomes important especially when considering that not all Wittgenstein's scholars are familiar with web-based educational tools. The five tabs can be briefly described as follows[5]:

1) the *Welcome* tab serves as a 'splash screen' and provides some contextual information and links to relevant resources;

2) the *Browse the text* tab (fig. 5) lets users browse the Tractatus, which is made available in three versions (the original German edition and the two major English translations). In order to facilitate this activity, a tree-like outline of the book on the left hand side lets them jump quickly to a specific paragraph. Moreover, we make use of a simple mechanisms for helping learners select which of the text's translations to visualize: when the mouse hovers one of the paragraphs shown on the right side of the screen, this is highlighted and a contextual menu appears above the text. By clicking on one of the available options, it is possible to view more than one translation at the same time;

3) the *Browse the annotations* tab supports a different type of text navigation by means of a smart-index of the topics associated to the Tractatus' fragments (fig. 6). For example, by clicking on a paragraph, it is easy to see all the the annotations which have been associated to it (in the *local* panel*)*. Similarly, by clicking on an annotation we can search for all the paragraphs related to it, which are displayed in the main central panel. Users can also go through all the philosophical annotations available (by means of the *categories* panel) or find out more information about an annotation in free-text form (*describe* panel) or by looking at what relations it entertains with other annotations (*inspect* panel).

[5] A lengthier description can be found in the first author's PhD thesis [29]

4) the *Browse the pathways* tab lets users select topics of interest and explore related resources by means of the 'learning pathways' facility (see the next section).

5) finally, the *Browse the ontology* tab visualizes the tree-hierarchy of the ontological representations PhiloSurfical relies on. This tab does not have any specific learning functionality but it has been added to the prototype mostly as a way for instructional technologists to inspect the underlying model of PhiloSurfical.

3.4. Ontology-enabled pathways for learning philosophy

The *learning pathways* are the most advanced navigation facility PhiloSurfical provides. By means of these pathways we aim at helping learners explore actively and autonomously the world of philosophy (see also section 1 above). In general, a 'pathway' is essentially a way to retrieve different instances stored in the knowledge base and organize them into a coherent whole. Pathways' results are normally presented in the form of *maps of connected ideas* - e.g., a map of competing views on the same topic, or a map of the philosophical problems typical of a research area - thus helping a student analyze a particular concept and interpret its significance within the various existing philosophical contexts.

So, for example, let us imagine a learning scenario where Lisa, a young philosopher, attempts to make sense of Wittgenstein's text. After having explored a number of Tractatus' topics by using the *Browse the annotations* tab described above (tab-3, fig. 6), Lisa develops more interest for the topic called "philosophy of mathematics". Thanks to the mechanisms available in tab-3, Lisa can see where this topic is dealt with by Wittgenstein in the text, and also how it is related to other topics. He/she then realizes that a key point to clarify concerns the significance of the so-called "problem of the foundations of mathematics". In order to benefit from more perspectives on this topic, our student now moves to the *Browse the pathways* tab. Here she can search and select the `problem` instance called "problem of the foundations of mathematics" and find out more about it by using the pathway called *problem-centric map of the attempts to solve a problem*. As shown in figure 7, this type of query produces a list of concurrent `view` instances which have been classified as attempting to solve that `problem`. Each `view` is presented together with other useful information too (e.g., the values of the slots *has-main-exponent*, *has-exemplar-theory*, etc.).

As a result, Lisa can now see what other authors have attempted to solve the "problem of the foundations of mathematics" - e.g., Plato and Frege. Also, she realizes that their respective theories have to be considered too when trying to understand this problem. In order to do so, she starts by selecting Frege's "mathematical logicism" instance and explore what pathways are available for it. Finally, in order to find some other literature about the topic, Lisa selects a *textual* type of pathway and arrives at the Stanford Encyclopedia of Philosophy entry about 'Frege's mathematical logicism'.

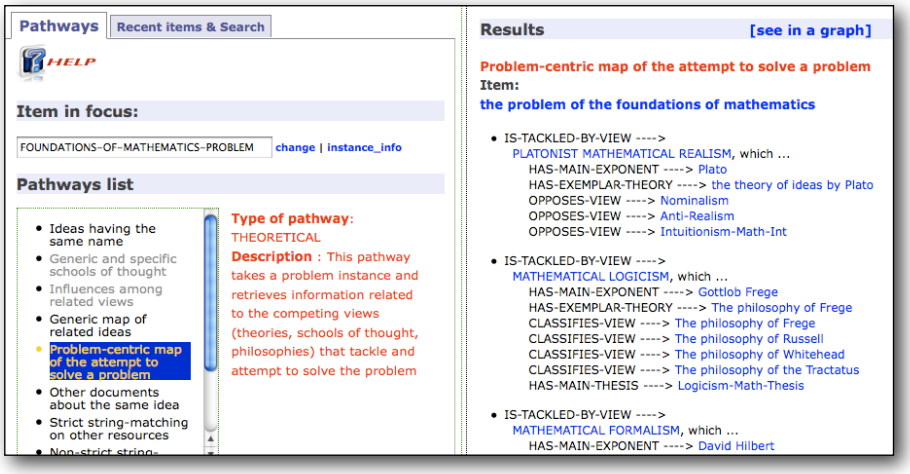

Figure 7. Pathway representing the various 'attempts to solve a problem'

At the interface level, such mechanisms can be described as follows. First of all, users select a content of interest as the starting point of a pathway (fig. 7, *item in focus* box). Learners may then choose from one of the available choices appearing in the *pathways list* panel (see figure 7, bottom-left). The pathways that are not available are dimmed out; the available ones, instead, come with a brief description explaining their meaning. Once triggered, the pathway's results are shown as a list of interrelated entities (fig. 7, *results* panel). Here, a number of important relations among the pathway's items are made explicit, so to highlight their significance in the philosophical discourse. Moreover, by clicking on any of these items it is possible to put it into focus and use it as the starting point of new pathways. A *recent items* panel is used to keep track of all the items selected since the beginning; also, from here it is possible to search for these topics elsewhere on the web (e.g., on philosophical portals, specialized search engines, etc.).

Furthermore, by clicking on the *see in a graph* button learners can view the pathways results' using a graphical visualization. E.g., in fig. 8 we can see the results of a theoretical pathway starting from the idea of "Frege's conception of logic". In this case the pathway selected is *generic map of related ideas*, which simply shows all information associated to an idea.

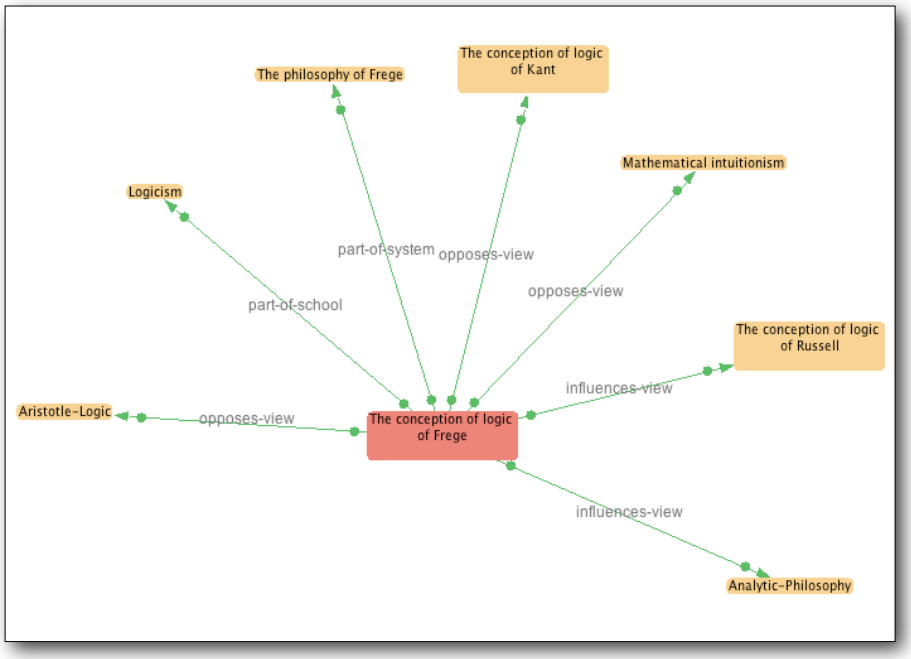

Figure 8. Graphical view of a theoretical pathway about Frege

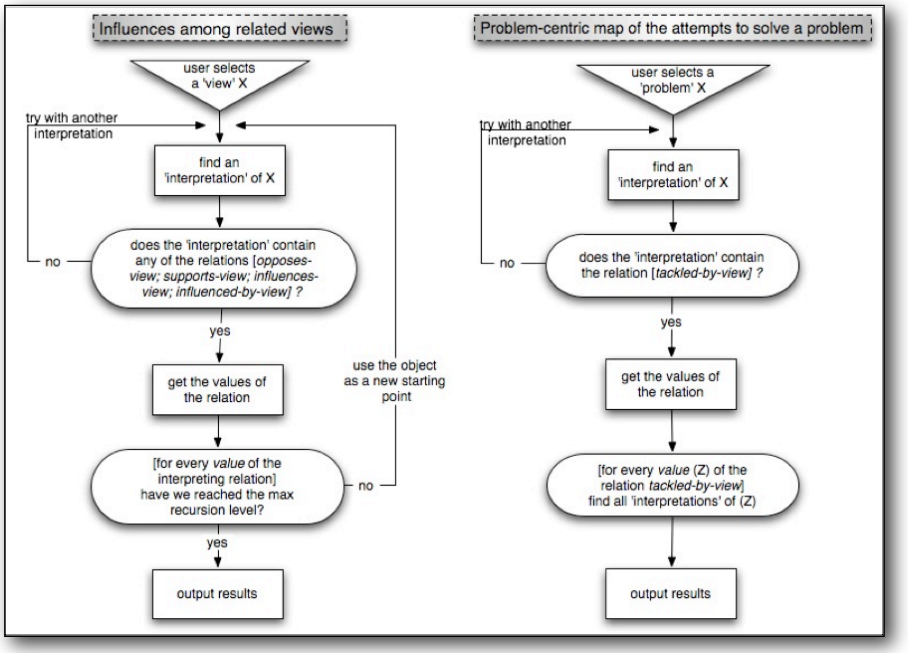

Figure 9. Abstract representation of two pathways' algorithm

We classified pathways according to the ontological type of their 'entry point', and, more generally, according to the types of the instances that are retrieved from the knowledge base. So, for example, by selecting instances of `philosophical-idea` we would usually trigger a *theoretical* pathway; instead, if we selected instances of `person` we would probably trigger a *textual* or *historical* pathway.

Because of space limitations, we cannot give here a complete description of all the pathways made available in PhiloSurfical[6]. In the table below (table 2) it is possible to see more information about a specific type of learning pathways, the *theoretical* ones.

Table 2. The theoretical pathways available in PhiloSurfical

Name (input type)	Description
Ideas having the same name (propositional-content)	This pathway retrieves ideas having the same name but a different meaning than the selected one. E.g., starting from the concept of 'fact' in Wittgenstein, we would find out about other authors who used the word 'fact' in a different sense (such as Frege and Russell).
"Generic and specific schools of thought" (school-of-thought)	Starting from a school of thought, this pathway retrieves a set of related schools of thought that are all *specializations* of the same generic one. This pathway is related to the formalization presented in section 3.5.4: e.g., by focusing on 'atomism' we would be able to see the related contextual versions of it, such as 'logical atomism', 'metaphysical atomism', 'social atomism', etc.
"Influences among related views" (view)	Starting from a view, this pathway is a recursive function showing information about other views that support/compete with the first one. E.g., starting from 'Wittgenstein's theory of language', we could go to the 'Russell's theory of language' (which opposes it), then to 'Whitehead's theory of logic' (which supports Russell's) etc.
"Generic map of related ideas" (propositional-content)	This pathway shows all the information an idea has been described with. This is a generic way to retrieve all the interpretations associated to an idea.
"Problem-centric map of the attempts to solve a problem" (problem)	This pathway takes a problem instance and retrieves information related to the competing views (theories, schools of thought, philosophies) that tackle that problem.

Finally, it is important to mention that internally PhiloSurfical represents pathways as abstract procedures applicable to any ontology-compliant data repository. For instance, in figure 9 we reproduced the algorithms behind the 'influences among related views' and the 'problem-centric map of the attempts to solve a problem' pathways (cf. also table 2 above). In particular, notice that after a pathway is triggered we normally scan the knowledge base for instances of the `interpretation` class mentioning the item which has been selected by the user. This class serves as an abstraction mechanisms for letting multiple annotators work together within PhiloSurfical; essentially, this means that every time an annotator formalizes a philosophical concept through the ontology, her activity is 'reified' by instantiating an interpretation object. A more detailed description of this feature can be found in another publication [21].

[6] A complete list of the pathways can be found in the first author's PhD thesis [29]

4. Related Work

In general, we reckon that there are two main major contributions in our work. First, an extensive ontology to represent the philosophical world (and in particular, philosophical ideas). Second, a Semantic Web oriented system for supporting learners in navigating interactively through philosophical resources. Accordingly, we will describe related research enterprises for that regards both the formal representation of philosophical domains and the navigation of them through semantic technologies.

With reference to the first aspect, the most relevant (and to our knowledge unique) attempt to *systematically* formalize the philosophical domain is the one carried out in [30], as part of a *digital library* project aimed at building a dynamic ontological-backbone for the online version of the Stanford Encyclopedia of Philosophy (SEP). Compared to our approach, this work is less focused on knowledge modeling and more targeted at finding useful information extraction techniques, which could benefit from the vast expert-reviewed SEP. For example, in their case the *idea* sub-branch of the ontology is populated according to "semantic relevance" of ideas (based on words co-occurrence), instead of trying to model a hierarchy of types. Therefore, we see the two approached as fundamentally complementary and likely to be used together in future work.

As various publications suggest, the *humanities computing* community has recently been more interested in the usage of ontologies for facilitating data representation and exchange [31,32]. In this context, the Discovery project [26] stands out for its explicit goal of creating an ontology-centered infrastructure usable by philosophers for exchanging data on the Semantic Web. In particular, the authors plan to use a "network of ontologies" [33]. This seems really promising, but unfortunately at the time of writing there is still no publicly available ontology for the philosophical domain. We plan to investigate how our results compare with theirs as soon as they will make them available.

Regarding the formalization of 'abstract' ideas (and especially philosophical ideas) we found little evidence of relevant work in the *knowledge representation* research literature. Although models such as Wordnet [34] and Cyc [35] have in their knowledge-base philosophy-related concepts, they present them in hierarchies that are either too flat (e.g. everything is a subclass of "doctrine") or not complex enough to support any navigation mechanism. When compared to such models, our ontology proved to be much more suited to the task[7]. Two noteworthy exceptions must be mentioned here. First, the DnS module of Dolce [36], which is "intended to provide a framework for representing contexts, methods, norms, theories, situations", and has strongly influenced us. However, our ontology appears to be much more specifically suited to represent philosophical entities, such as schools of thoughts or problems. In fact, such topics are only marginally treated by DnS, which focuses on the formalization of entities such as plans, laws and regulations (legal objects).

Second, the research of Mizoguchi and colleagues. His ontology of 'representations' [37] includes a conceptual model which organizes propositional contents in two groups, *product* propositions and *design* propositions. The former "works as specification of the production of something", while the latter "is the product". We have found this distinction very useful and included it in our formalizations. Also, our modeling of philosophical theories (cf. section 2.3) can be compared to their formalization of learning theories in the OMNIBUS ontology [38].

[7] The reader can find a more detailed analysis and comparison of the philosophical concepts in CYC and other foundational ontologies in chapter 3 of the first author's PhD thesis [29]

The authors present a "theory-neutral" ontology that aims at expressing the similarities and differences of various instructional and learning theories. Their approach is based on the "working hypothesis that a sharable engineering approximation related to learning can be found in terms of the changes that are taking place in the state of the learners". Consequently, the authors' characterization of learning theories relies on an extensive descriptions of learners' states and other important contextual elements of learning scenarios. In general, this approach seems to be an interesting alternative to ours. In fact, we only attempted to model theories or schools of thought according to their 'theoretical' features, i.e., without referring to their implications in the real world (e.g., the change a learners' state). This might have been a direct consequence of the fact that often philosophical theories have a much less 'pragmatic' attitude, especially when compared to learning and instructional theories. However, we think that this problem necessitates further research so we plan to investigate it in future work.

Finally, our formalization of fields of studies (cf. section 2.4) could be related to the various work done in *digital libraries* subjects' classification. Although we come from a different perspective, we acknowledge that approaches such as the *mereotopological* one [39] could be well suited also for the philosophical domain.

The second contribution in our work regards the *semantic navigation* component of the PhiloSurfical tool. In this respect, the most relevant research work it could be compared to is Story Fountain[40]. This is an ontology-based application developed to support a community in the exploration of digital resources, specifically *stories*. Users ask questions about the domain (Bletchley Park, a second-world-war heritage site) and receive as answers some explicatory *paths* along the many annotated stories in the knowledge base. Our pathway-centered approach have been largely inspired by Story Fountain, although our application domain - philosophy - required a radical change of perspective. In fact, while Mulholland and colleagues are creating pathways that focus on stories' protagonists (e.g., an army colonel) and objects (e.g., a pistol), in our scenario those type of entities are often secondary. The paths we are dealing with usually center around abstract ideas, such as philosophical theories and problems.

Finally, it is worth mentioning recent research aimed at facilitating the semantic navigation of digital resources' repositories, for it *complements* our learning-pathways approach. *Faceted browsing* systems usually provide generic architectures that aim at letting users explore potentially unfamiliar domains in a gradual and incremental manner. These approaches, inspired by faceted theory [41], have been tested in various humanities domains, such as classical music [42], visual arts [43], cultural heritage [44] and literature [45]. In general, by means of highly interactive visualization mechanisms which are controlled by the user's selection of facets, the structure of a domain can be disclosed in a very intuitive manner. The main limitations of these systems, in our opinion, is linked to their very best feature. That is, being largely non-domain specific and allowing navigation based on 'small' and 'incremental' steps (i.e. selection of views/facets) the navigation mechanisms can hardly be tailored to specific learners' needs. For instance, it would not be possible to construct a 'view' which organizes resources in a way that mimics, or at least supports, the traditional ways a discipline is presented or taught. In conclusion, our narrative inspired approach seems to be better targeted to an educational scenario.

5. Conclusions

In this chapter we summarized our work with the PhiloSurfical tool. This is an application built to support students in understanding a philosophical text, through contextual navigation mechanisms based on semantic technologies. The application is being prototyped with Wittgenstein's *Tractatus Logico-Philosophicus* using a philosophical ontology we created and instantiated with the relevant data. The ontology modeling process has demonstrated to be crucial to the aim of providing valuable and non-naïve navigation mechanisms. In particular, we showed how the usage of solid modeling schemas can serve to solve ambiguities in the philosophical domain, and possibly to tidy up poorly or wrongly structured data in the quickly improving Semantic Web. We are currently in the process of elaborating the data obtained from two separate evaluations, one of the application and one of the ontology. We plan to make such results available in a separate publication.

Acknowledgments

This work has been carried out under a grant provided by the EU-funded Knowledge Web project. We would like to thank all the people who have been providing feedback and support during the various stages of the research. In particular (in chronological order) Andrea Bernardi, Keith Frankish, Gordon Rugg, Marian Petre, Riichiro Mizoguchi and Martin Doerr.

References

[1] E. Duval, Learning Technology Standardization: Making Sense of it All. *International Journal on Computer Science and Information Systems* **1** (2004), 33-43.

[2] D. Gasevic, J. Jovanović and V. Devedzic, Enhancing Learning Object Content on the Semantic Web. IEEE International Conference on Advanced Learning Technologies (ICALT'04), 2004.

[3] International Workshop on Applications of Semantic Web technologies for E-Learning (SW-EL), official website: http://compsci.wssu.edu/iis/swel/index.html, retrieved on July 2009.

[4] T. R. Gruber, A Translation Approach to Portable Ontology Specifications. *Knowledge Acquisition* **5** (1993), 199-220.

[5] D. Gasevic, J. Jovanović, V. Devedžić and M. Bošković, Ontologies for Reusing Learning Object Content. International Workshop on Applications of Semantic Web Technologies for E-Learning (SW-EL), ICALT-05 conference, 2005.

[6] J. Brase and W. Nejdl, Ontologies and metadata for eLearning, in *Handbook on Ontologies* (eds S. Staab & R. Studer), Springer-Verlag, 2004, 555-574.

[7] L. Stojanovic, S. Staab and R. Studer, Elearning Based on the Semantic Web. WebNet2001, 2001.

[8] K. M. Brooks, Do Story Agents Use Rocking Chairs? The Theory and Implementation of One Model for Computational Narrative. Fourth ACM international conference on Multimedia, 1997.

[9] A. Gangemi, Ontology Design Patterns for Semantic Web Content, International Semantic Web Conference , ISWC'05, 2005.

[10] L. Wittgenstein, *Tractatus Logico-Philosophicus*, Routledge & Kegan Paul, 1922.

[11] J. Bruner, *The Process of Education*, Harvard University Press, Cambridge, MA, 1960.

[12] J.S. Brown, A. Collins and P. Duguid, Situated Cognition and the Culture of Learning, *Educational Researcher* **18** (1989), 32-42.

[13] D. Laurillard, *Rethinking University Teaching*, Routledge, 1993.

[14] G. Kemerling, Teaching Philosophy on the Internet, 20th World Congress of Philosophy, 1998.

[15] W. Mays, The Teaching of Philosophy, 13th University Conference, 1965.

[16] T. Kasachkoff, *Teaching Philosophy: Theoretical Reflections and Practical Suggestions*, Littlefield Publishers, 2004.

[17] A. Carusi, Taking Philosophical Dialogue Online. *Discourse: Learning and Teaching in Philosophical and Religious Studies* **3** (2003), 95-156.
[18] M. Doerr, The Cidoc Conceptual Reference Module: An Ontological Approach to Semantic Interoperability of Metadata. *AI Magazine archive* **24** (2003), 75-92.
[19] N. Crofts, M. Doerr, T. Gill, S. Stead and M. Stiff, Cidoc-CRM Version 4.2 - Reference Document, 2005.
[20] E. Motta, *Reusable Components for Knowledge Modelling - Principles and Case Studies in Parametric Design Problem Solving*, IOS Press, The Netherlands, 1999.
[21] M. Pasin and E. Motta, Ontological Requirements for Annotation and Navigation of Philosophical Resources. To appear in *Synthese Special Issue: Representing Philosophy*, 2009.
[22] M. Crampes, S. Ranwez, Ontology-Supported and Ontology-Driven Conceptual Navigation on the World Wide Web, 11th ACM Hypertext Conference, 2000.
[23] S. Chatman, *Story and Discourse*, Cornell University Press, 1978.
[24] M. Pasin and E. Motta, Semantic Learning Narratives, International Workshop on Applications of Semantic Web Technologies for E-Learning (SW-EL), KCAP'05, 2005.
[25] S. Auer, C. Bizer, G. Kobilarov, J. Lehmann, R. Cyganiak, and Z. Ives, DBpedia: A Nucleus for a Web of Open Data, 6th International Semantic Web Conference (ISWC'07), 2007.
[26] Discovery project, official website: http://www.discovery-project.eu/home.html, retrieved on July 2009.
[27] Lispworks - The integrated cross-platform development tool for ANSI common lisp, http://www.lispworks.com/, retrieved on July 2008.
[28] E. Weitz, HUNCHENTOOT - The Common Lisp web server formerly known as TBNL, http://weitz.de/hunchentoot/, retrieved on July 2008.
[29] M. Pasin, Ontological Requirements for Supporting Smart Navigation of Philosophical Resources. Ph.D. Thesis, Knowledge Media Institute, The Open University, 2009.
[30] M. Niepert, C. Buckner and C. Allen, A Dynamic Ontology for a Dynamic Reference Work, Joint Conference on Digital Libraries (JDCL'07), 2007.
[31] G. Nagypál, R. Deswart and J. Oosthoek, Applying the Semantic Web: The Vicodi Experience in Creating Visual Contextualization. *Literary and Linguisting Computing* **20** (2005), 327-349.
[32] J. M. Vieira and A. Ciula, Implementing an RDF/OWL Ontology on Henry the Iii Fine Rolls, OWLED, ESWC'07, 2007.
[33] M. Nucci, S. David, D. Hahn and M. Barbera, Talia: A Framework for Philosophy Scholars, 4th Italian Semantic Web Workshop (SWAP'07), 2007.
[34] C. Fellbaum, (ed.) *WordNet: An Electronic Lexical Database*, MIT Press, 1998.
[35] D. B. Lenat and R.V. Guha, *Building Large Knowledge-based Systems: Representation and Inference in the Cyc Project*, Addison-Wesley, Boston, Massachussets, 1990.
[36] A. Gangemi and P. Mika, Understanding the Semantic Web Through Descriptions and Situations, International Conference on Ontologies, Databases and Applications of Semantics (ODBASE), 2003.
[37] R. Mizoguchi, Tutorial on Ontological Engineering - Part 3: Advanced Course of Ontological Engineering, *New Generation Computing* **22** (2004), 198-220.
[38] Y. Hayashi, J. Bourdeau and R. Mizoguchi, Using Ontological Engineering to Organize Learning/Instructional Theories and Build a Theory-Aware Authoring System, *International Journal of Artificial Intelligence in Education* **18** (2008).
[39] C. Welty and J. Jenkins, Formal Ontology for Subject. *Journal of Data and Knowledge Engineering* **31** (1999), 155-181.
[40] P. Mulholland, T. Collins and Z. Zdrahal, Story Fountain: Intelligent Support for Story Research and Exploration, 9th International Conference on Intelligent User Interface, 2004.
[41] S.R. Ranganathan, *Elements of Library Classification*, South Asia Books, 1990.
[42] mc. Schraefel, D. Smith, A. Russel, A. Owens, C. Harris and M. Wilson, The Mspace Classical Music Explorer: Improving Access to Classical Music for Real People, V MusicNetwork Open Workshop: Integration of Music in Multimedia Applications, 2005.
[43] M. Hildebrand, J. Van Ossenbruggen and L. Hardman, /Facet: A Browser for Heterogeneous Semantic Web Repositories, International Semantic Web Conference (ISWC'06), 2006.
[44] E. Hyvönen, T. Ruotsalo, T. Haggstrom, M. Salminen, M. Junnila, M. Virkkila, M. Haaramo, T. Kauppinen, E. Makela, K. Viljanen, CultureSampo-Finnish Culture on the Semantic Web. The vision and first results, *Information Technology for the Virtual Museum*, LIT Verlag, 2008.
[45] B. Nowviskie, Collex: Facets, Folksonomy, and Fashioning the Remixable web, Digital Humanities conference (DH'07), 2007.