

Ontological requirements for annotation and navigation of philosophical resources

Michele Pasin · Enrico Motta

Received: 6 July 2008 / Accepted: 22 January 2009 / Published online: 29 September 2009
© Springer Science+Business Media B.V. 2009

Abstract In this article, we describe an ontology aimed at the representation of the relevant entities and relations in the philosophical world. We will guide the reader through our modeling choices, so to highlight the ontology's practical purpose: to enable an annotation of philosophical resources which is capable of supporting pedagogical navigation mechanisms. The ontology covers all the aspects of philosophy, thus including characterizations of entities such as people, events, documents, and ideas. In particular, here we will present a detailed exposition of the entities belonging to the *idea* branch of the ontology, for they have a crucial role in the world of philosophy. Moreover, as an example of the type of applications made possible by the ontology we will introduce PhiloSurfical, a prototype tool we created to navigate *contextually* a classic work in twentieth century philosophy, Wittgenstein's Tractatus Logico-Philosophicus. We discuss the potential usage of such navigation mechanisms in educational and scholarly contexts, which aim to enhance the learning process through the serendipitous discovery of relevant resources.

Keywords Ontology · Philosophy · Digital narratives · Knowledge representation · Semantic web · CIDOC · Tractatus Logico-Philosophicus

1 Introduction

The Semantic Web augments the web with a layer of data, called metadata, which formally describes information. The idea here is to develop a large-scale repository of

M. Pasin (✉) · E. Motta
Knowledge Media Institute, Open University, Milton Keynes, UK
e-mail: m.pasin@open.ac.uk

E. Motta
e-mail: e.motta@open.ac.uk

formally characterized resources, over which intelligent agents would perform various kinds of operations for the user (Berners-Lee et al. 2001). As part of this effort, our research focuses on the definition of the appropriate metadata which could be used to describe philosophical resources. In particular, the approach we are taking is further characterized by the fact that we want to make use of such metadata in a *pedagogical* scenario. This set of descriptors, codified in an *ontology* (a formalized conceptual organization to support the encoding of metadata) (Noy and McGuinness 2001), can then be used to provide intelligent mechanisms for selecting and navigating through learning materials. Moreover, by linking metadata to relevant explanatory and exegetical materials, we will give students additional means for contextualizing philosophical resources.

In general terms, the pedagogical principle inspiring us reflects the idea of an ‘invisible guide’, able to support ‘smart’ navigation by discovering interesting connections between metadata and philosophical resources. This approach can be implemented in three steps. First, we let experts (e.g. philosophy teachers) *represent* part of their knowledge using the ontology—i.e. they *instantiate* the ontology using contents related to a philosophical subject of choice. For instance, in the context of annotating Wittgenstein’s Tractatus, an expert may instantiate the generic notion of *school-of-thought* with the concept of “logical atomism”. Second, our experts *annotate* learning resources using the metadata just created—i.e. they formally associate one or more instances to a learning resource. For example, they may associate “logical atomism” with a specific statement in the Tractatus. Third, we construct algorithms which, by drawing on the ontological categories and the experts’ annotations, can organize dynamically the presentation of learning resources. For example, by bringing in other resources related to schools of thought that *oppose* “logical atomism”. In other words, resources can be viewed according to a specific perspective, which can be *historical, theoretical, geographical*, etc. This results in a series of navigation mechanisms for students to *explore* such resources in an unsupervised manner. In a nutshell, the ontology-based annotation would bring ‘authoritative structure’ to learners’ autonomous explorative activities.

In this scenario, the ontology is similar to an invisible map that helps students moving through learning resources by means of pre-defined *learning pathways*. As discussed elsewhere (Mulholland et al. 2004), it is important to remember that the underlying assumption of this approach is that the ontology-based system is not supposed to provide a specific answer to the questions a learner or researcher may pose to it; instead, its goal is to facilitate the discovery of related (and possibly unknown) resources where the answer can be found. This is achieved by making transparent a number of *coherence principles* typical of the philosophical discourse (e.g. a *historical* evolution of a school of thought, the *theoretical* implications of an argument, etc.).

Given these premises, we can describe the research work presented in this article as an attempt to construct a formal *meta-language* that allows the categorization of philosophical subjects. In defining this meta-language, we have taken inspiration from concepts that are commonly used for characterizing philosophical scholarship. For example, we included in the ontology notions such as *philosophical system, argument* and *school of thought*. This does not mean that we are prescribing a particular

usage of these concepts; different authors can in fact characterize the same philosophy in different ways.

This important aspect may be clarified through an example. Our ontology strictly defines a *philosophical-system* as a type of object which can have the property *part-of-school-of-thought* (cf. Sect. 3.5.3), but it does not specify any *philosophical-system* in particular. Annotators and domain experts are expected to do this—and quite surely, they will do it with a great degree of subjectivity. For example, in a Wittgensteinian context some will consider “the philosophy of the first Wittgenstein” and “the philosophy of the second Wittgenstein” as separate instances of *philosophical-system*, while others could find the distinction quite unreasonable and instead define a single instance of *philosophical-system*, which represents “Wittgenstein’s philosophy” as a whole.

What we are pointing out here is that the ontology supports *both interpretations*. This is a consequence of the fact that we gave our representations (classes and relations) a high level of generality i.e. we wanted them to be as re-usable as possible, especially among annotators having different philosophical views. The main advantage of this approach is that even when two authors’ interpretations are radically different, if the underlying meta-language is the same, we can still create connections between the alternative models.

Of course, we are not claiming that there can be only one ontology for this purpose. And indeed our work has been driven by a very specific objective, i.e. the creation of navigation mechanisms which are pedagogically interesting and computationally feasible. Other philosophical ontologies with different modeling choices and rationales are likely to be created in the future. In such cases, ontology-mapping techniques (Kalfoglou and Schorlemmer 2003) could be investigated so to guarantee interoperability among heterogeneous models.

The rest of the article is organized as follows: the next paragraph (1.1) gives a few technical notes which will facilitate readers in understanding the rest of this article. Section 2 summarizes a number of ontological requirements and the generic approach we used to satisfy them. Section 3 focuses on the description of the classes and relations representing *philosophical ideas*. Section 4 introduces PhiloSurfical, a prototype tool exemplifying the use of ontology-based navigation mechanisms within a pedagogical scenario. Finally, Sect. 5 contains some references to related projects.

1.1 Technical notes

From the implementation point of view, the ontology (which at the time of writing counts 348 classes) is formalized by using the Operational Conceptual Modelling Language (OCML) (Motta 1999), which provides rich support for both specification and reference. Import/export mechanisms from OCML to other languages, such as OWL (W3C 2004) and Ontolingua (Farquhar et al. 1996), ensure interoperability with knowledge representation standards.¹

¹ The latest version of the ontology can be found online at <http://www.philosurfical.open.ac.uk/ontology/>.

In the rest of the article, when examples from the ontology are provided, we use the OCML syntax for describing classes, instances and rules. In order to facilitate the reading of this article 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”). In the figures, classes are oval-shaped, rounded rectangles stand for instances and arrows represent relations. In particular, if not labelled otherwise, dashed arrows stand for the *instance-of* relation, while solid arrows stand for the *subclass-of* relation.

As a final remark, we invite the reader who is not familiar with the *knowledge representation* approach and terminology to consult the relevant literature, since such an understanding is crucial in order to fully comprehend our work. In fact, we must remember that although ontologies have their roots in philosophy, their computational equivalents have raised a number of research problems which were previously unseen in philosophy (Zúñiga 2001). Unfortunately, a discussion of such issues would exceed the scope of this article. Readers may find a good introduction to the topic in the *course on ontological engineering* by Riichiro Mizoguchi (Mizoguchi 2004).

2 Philosophy as a domain for knowledge representation

2.1 Domain analysis

In order to identify an initial set of ontology requirements, we used various *informal* knowledge acquisition techniques.² Mainly they consisted of discussions with domain experts, analyses of the implicit curricula formalized in philosophical textbooks, consultation of traditional encyclopedias and online philosophy directories. Then we also carried out a more *formal* knowledge acquisition experiment:³ a group of domain experts (lecturers and Ph.D. students) were involved in a card-sorting task (Rugg and McGeorge 2005) aimed at identifying some mechanisms practitioners employ for classifying philosophical entities (especially *abstract* entities, i.e. ideas).

In general, these results led us to conclude that a suitable semantic model should provide support for representing:

- (a) *historical events*, that is, events which are inherently time-dependent (e.g. the publication of a book, or an author’s subscription to a viewpoint);
- (b) generic *uncertainty*, since often we are talking about facts which cannot be located exactly in the time and space dimension (e.g. the birth of Heraclitus);
- (c) *information objects*, and especially language-based information objects, as they are the traditionally preferred medium philosophical contents are expressed with;
- (d) *interpretation events*, intended as the process of attributing an abstract content to an information object (e.g. when we say that ‘Aristotle’s fourth book of the *Metaphysics* states an ontological principle’);

² A good overview of the various techniques available can be found at <http://www.epistemics.co.uk/Notes/63-0-0.htm>

³ We intend to make available all the results of the experiment in the near future, in a separate publication.

- (e) coexistence of *contradictory information*, which is a direct consequence of D (e.g. when people claim different or opposing views on the same subject);
- (f) *viewpoints*, and other non-material entities (‘philosophical ideas’), for they are the objects philosophers are usually involved with, in their everyday practices;
- (g) *varying granularity*: this feature refers to the fact that philosophers normally (re)define the questions and ideas which lie at the centre of their work. As a result, the conceptions of two philosophers can have very little in common, if not at a meta-level. Thus, our model should be capable of overcoming the difficulties imposed by the ‘radicalism’ of philosophical conceptions. This means providing facilities to properly describe a philosophy, considered in both its singularity and within an historical perspective. For example, being able to express the historical contiguity of ‘Aristotle’s distinction of the four causes’ with ‘Hume’s radical skepticism regarding the cause notion’, although the two conceptions, taken singularly, do not have much in common with respect to the definition of the ‘cause’ notion.

2.2 Overview of the ontology

The main feature characterizing our ontological approach is the decision to employ the CIDOC Conceptual Reference Model (Doerr 2003) as a starting point for our formalizations. The CIDOC-CRM ontology started out as an attempt of the committee of the International Council of Museums (ICOM) to achieve semantic interoperability for museum data. Since 1996, the formal model has improved considerably till becoming in 2006 an ISO standard (Crofts et al. 2005). It is now (version 4.2) in a very stable form, and contains 75 classes and 108 properties, both arranged in multiple *is-a* hierarchies. The choice of using CIDOC-CRM was motivated by two reasons.

First, because of its widely recognized status as a *standard for modeling cultural heritage data*. In fact, by reusing and extending an existing and internationally recognized ontology, we will give our users more chances to benefit from the emerging Semantic Web infrastructure.

Second, 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, we believe this cannot be examined without an adequate consideration of the *historical* dimension, that is, a history of the events related (directly or indirectly) to those ideas. Thus, with reference to the domain analysis described above, we can say that point A is directly addressed by CIDOC’s generic modeling approach.

As an example, in Fig. 1, we can see a typical event-centered instantiation of 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”).

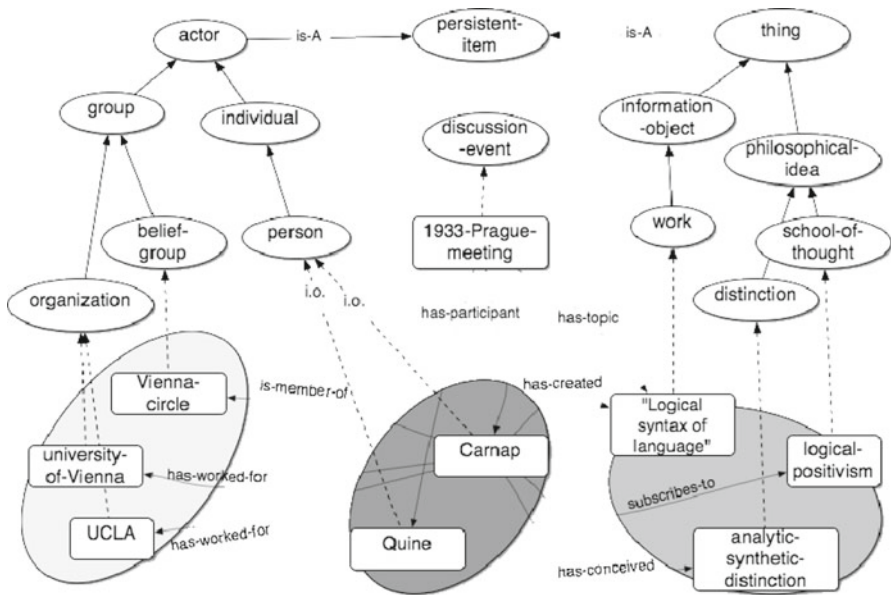


Fig. 1 A typical event-based modeling in CIDOC

This kind of modeling, in the context of the PhiloSurgical tool (cf. Sect. 4), is extremely useful because of the multiple navigational pathways it can support (e.g. imagine a ‘lateral’ step taking us 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.).

As already pointed out in previous publications (Pasin and Motta 2007; Pasin et al. 2007), we decided to integrate the event-based CIDOC reference model with formalizations from other ontologies, because they provide facilities that are relevant to the points C, D and E we have highlighted earlier during the domain analysis. For example, we included a time-reasoning library implementing the well-known Allen specifications (Allen 1984); we included knowledge about the domain of publications from the AKT reference ontology (AKT 2002) and knowledge about information objects from the related module of the Dolce foundational ontology (Gangemi et al. 2002).

Finally, a large portion of the ontology consists of new concepts and relations, mostly aimed at the description of philosophical *events* and *ideas*. The events having more relevance with respect to the philosophical domain are the following:

- (1) the temporal entities regarding *events related to the academic life and to the life of philosophers*. Among this group of events, we have *births* and *deaths* of philosophers (e.g. the death of Socrates), *production* of physical objects (e.g. Pascal’s construction of the arithmetic machine), *journeys* performed during their lives (e.g. Wittgenstein’s trip to Norway), *production* of publications (e.g. the publication of the first English version of Kant’s “Critique of practical

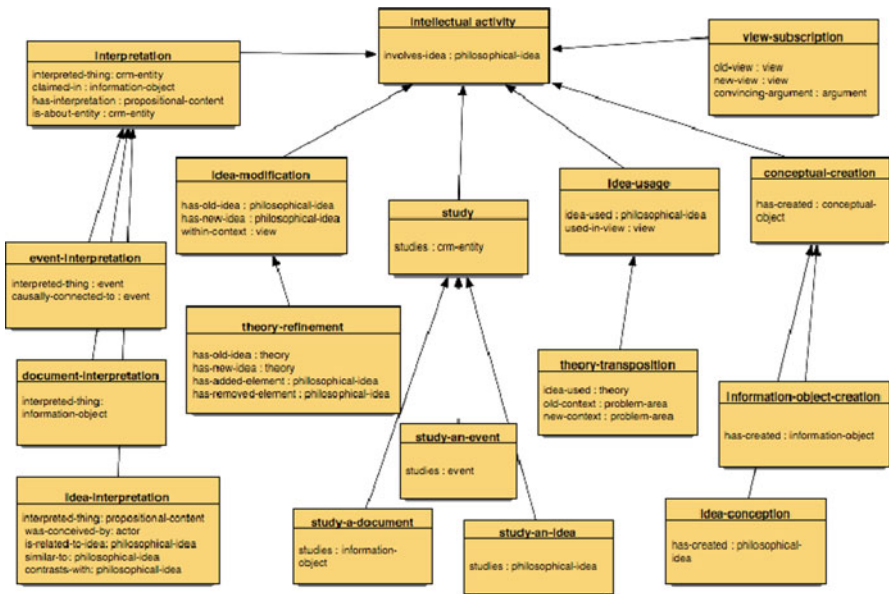


Fig. 2 The intellectual activities branch

reason” in 1836), social-gatherings, conferences, joining of groups (e.g. when Aristotle joined the Academy of Plato, or when Heidegger joined the Nazi party).

- (2) The temporal entities related to the *production and modification of philosophical ideas*. These types of events are gathered under the class intellectual-activity (see Fig. 2). Among them, we can find conceptual-creation (event modeling the creation of conceptual entities such as ideas and information objects); idea-modification (events reflecting the changing of one or more ideas within the context of a view, e.g. the evolution of the meaning of “libido” in the work of Sigmund Freud); theory-transposition (class modeling the special case when a theory is taken out of a context and reused within another one, e.g. “Spencer’s evolutionism”, which extends “Darwin’s evolutionism” from biology to metaphysics), etc.
- (3) The temporal entities representing *philosophical historical periods*, i.e. macro-events (in CIDOC, such entities are subsumed by a class named period) characterized by an intrinsic reference to a specific group of people or a school of thought. The important classes here are intellectual movement (e.g. the “enlightenment”) and philosophical movement (e.g. “logical positivism”, interpreted as an event). The formal framework used for representing the characteristics of these entities has been previously discussed (Pasin and Motta 2007) under the title ‘pattern #1: is rationalism a school of thought or an event?’

The other major section which we extended CIDOC with is the one departing from the philosophical-idea class, which is located in the conceptual-object branch of the ontology (according to CIDOC, this is where all abstract entities are).

In relation to the initial domain analysis, these formalizations satisfy the requirements described in points F and G. In Sect. 3, we concentrate the discussion on this branch of the ontology.

2.3 Support for alternative interpretations

It is important to remember that only classes and relations are what remains unchangeable in our system, i.e. that is where lies the *ontological commitment*⁴ we demand from anybody using the ontology. On the contrary, the instantiation of our classes with elements specific to a single philosophy is a process which relies entirely on a user's private understanding of that philosophy.

As already discussed in Sect. 1, this feature allows annotators to use our meta-language with a great degree of freedom. As a result, the interpretations of philosophical subjects they create can be very different from each other. The only downside, in such cases, is that the results of incompatible instantiations are not handled easily by computers, thus requesting a manual integration.

In order to provide a solution to this problem, the ontology features a mechanism by which we can construct alternative and possibly competing interpretations of the same entity, in such a way that the computer 'knows' how to handle each interpretation as an alternative view on a common topic. This mechanism becomes useful, for example, when we want to have multiple annotators working *simultaneously* within a single ontology-based environment (i.e. because we are interested in highlighting with precision how the various people's interpretations differ).

The *interpretation* class, a subtype of *event*, is meant to abstract the act of *interpreting* something, intended as the process by which we attribute a meaning to an object (cf. also Fig. 3). In ontological terms, this translates to associating an instance of *propositional-content* (i.e. the idea representing the *interpretation*) to any other instance of the ontology (i.e. the *interpreted-thing*). Of course, since an *interpretation* is also an event, it inherits various properties which capture useful information such as the author of the interpretation (*carried-out-by* property), the time it was made (*has-time-specification* property), etc. For example, in our Wittgenstein-related knowledge base, we can have the following instance (see also Sect. 4.1):

```
(def-instance interpretation-001 concept-
interpretation
((interprets a-posteriori-by-wittgenstein)
 (has-related-concept experience-by-wittgenstein
  prop-of-science-concept)
 (has-opposite-concept a-priori-by-wittgenstein
  laws-of-logic-concept)
 (is-equivalent-to form-prop-science-concept))
```

⁴ Ontological commitments can be defined as "agreements to use the shared vocabulary in a coherent and consistent manner" (Gruber 1993).

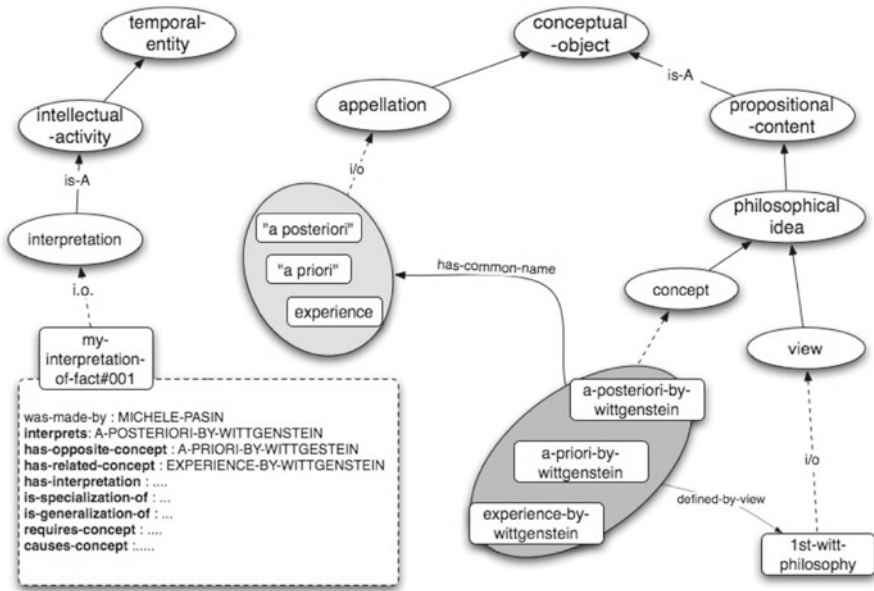


Fig. 3 Reification of ideas through the interpretation class

```
(is-related-to-idea mesh-metaphor fate-science-analogy)
(carried-out-by michele-pasin)
(has-time-specification 10-dec-08))
```

In this case, we are describing the properties of the concept of “a-posteriori by Wittgenstein” in such a way that these descriptions will be associated only to a specific user (i.e. the value of the property *carried-out-by*) and time (i.e. the value of the property *has-time-specification*). This is possible because the *concept-interpretation* class (a further specification of *idea-interpretation*) possesses all the properties normally used for describing an instance of *concept*: for example, *has-related-concept*, *has-opposite-concept*, *is-equivalent-to*, etc. (cf. Sect. 3.7). The result of this ‘reification’ mechanism is that we can have different descriptions of the same concept (and, in general, of any idea⁵) coexisting within the same knowledge base. In other terms, we are providing support for *concurring* and possibly *contradictory* information management. In future versions of our work, this feature is likely to be further developed with more a complex mechanism to retrieve, for example, contrasting interpretations, or letting users navigate through alternative views of the same ideas.

Finally, as represented graphically in Fig. 3, notice how the ontology allows also separating the *name* of an idea (through the *appellation* class) from the idea itself (an instance of *philosophical-idea*). For space reasons, we will not describe

⁵ In fact, by using a similar approach we created also other subtypes of *interpretation*, so to match all the remaining subtypes of *philosophical-idea*.

this feature here, but let us underline that it is a mechanism provided by CIDOC-CRM. In particular, within our philosophical context, this can be useful for describing multiple *linguistic translations* of the same idea.

3 A formal model for describing philosophical ideas

The ontological approach presented in the previous sections accounts mainly for the *factual* and *temporal* dimensions of the philosophical domain (e.g. we modelled entities such as *people*, *events* or *documents*). We must now investigate another dimension that is eminently philosophical, i.e. the *theoretical* one, the realm of philosophical *ideas*.

Where should one start in order to formalize the *types* of abstract entities discussed in philosophy? This seems a really puzzling question, and probably totally nonsensical to some. Such a slippery and debated domain, in fact, appears to challenge any stable formalization, and defeat any meaning-agreement process.

On the other hand, modern day digital phenomena such as the incredible growth of available information or the increasing need for interoperability standards call for a solution which, although inevitably partial and non-definitive, can bring many more advantages than no solution at all. As claimed by the authors of a recent project for the indexing of the Stanford Encyclopedia of Philosophy (Niepert et al. 2007):

“while no single ontology can possibly capture the full richness and interrelatedness of philosophical ideas, we are operating on the principle that having (at least) one ontology is better than none.”

In the light of this simple but important reflection, we have attempted to model commonly used philosophical concepts without taking any particular philosophical position, that is, for what is possible, trying to remain “outside” specific philosophical stances. Not doing so would have caused a multiplication of ontologies and definitions, each of them reflecting the world according to a single thinker.

Our approach, which can be related to a *constructivist* epistemology (Bachelard 1938), sees every philosophy as a system of interrelated conceptual entities which make sense of the world. From this perspective, we can say that such entities are all abstract (non-physical), since they are ‘what we use’ to refer to the physical world. The main consequence of this perspective is that even a common concept like “fire”, which would be normally instantiated as a physical entity, in our model becomes an instance of a concept (which is possibly related to a physical entity). In fact, the notion of fire, as any other notion, is socially constructed (Vygotsky 1978) and often explicitly defined by a viewpoint (e.g. the “Newtonian physics”, or the “philosophy of Heraclitus”). The fact that a generic agent happens to be more or less explicitly aware of this viewpoint, in all its aspects and subtleties, constitutes another issue and does not disprove the existence of it.

For us, the problem to tackle is therefore the *individuation of the types of non-physical-objects playing a role in the construction of viewpoints*, and, more broadly, having a recognizable function in the process of interaction and succession of viewpoints within the whole history of thought. As previously discussed (cf. Sect. 1),

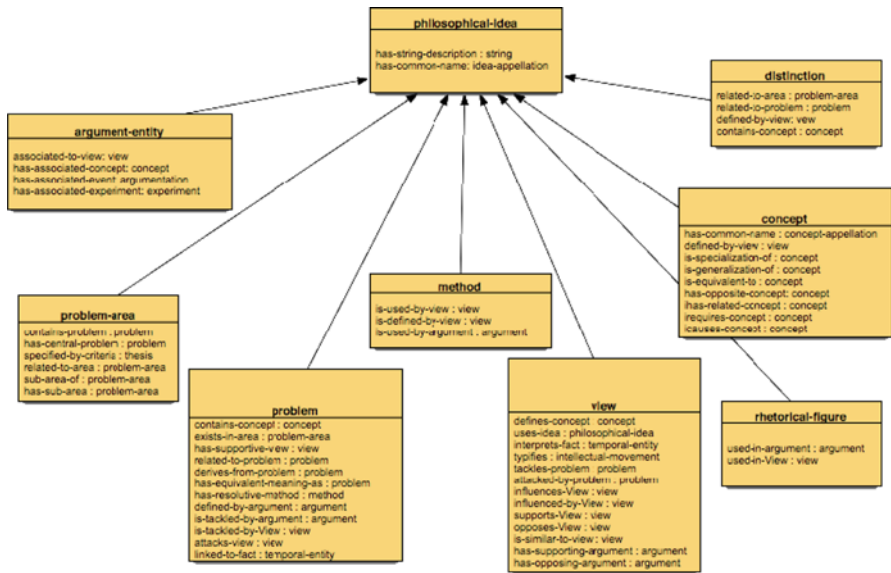


Fig. 4 The main classes of the philosophical-idea branch

the pragmatic requirements of creating a model which is at the same time *vastly reusable* and useful for creating *pedagogical learning pathways* have driven much of the formalizations presented below.

In total, we identified eight main types of philosophical ideas (see Fig. 4). The following sections discuss them in greater details.

3.1 Argument-entity

With the `argument-entity` class, we decided to group together two sets of related classes: `argument` and `argument-part` (see Fig. 5).

The first one is the reification of the `argumentation` class (which is a subtype of `event`), as it ‘freezes’ an actual argumentation between two or more thinkers into an abstract idea (i.e. an entity outside space and time). In previous versions of the ontology, we also named it ‘argumentative-knot’. In fact, it refers to famous focal points of the philosophical argumentation, where all the main argumentative threads converge and meet. These knots usually originate with one author, and subsequently recalled and reused (maybe in different domains or for different purposes) by other authors. So, for example, we can have the “third-man argument” of Plato, the Cartesian “*cogito-ergo-sum*” or the Kantian “transcendental deduction”. An important property of this class is *uses-method*, whose range is `argumentative-method` (a subclass of `abstract-method`), because through it we can specify, for example, a *deductive-argument*, an *inductive-argument* or an *abductive-argument*.

The second subclass of `argument-entity` is instead `argument-part`, which precisely serves to map out all the argumentative steps of a standpoint. For the moment, we only defined *assumption*, *demonstration*, *conclusion* and

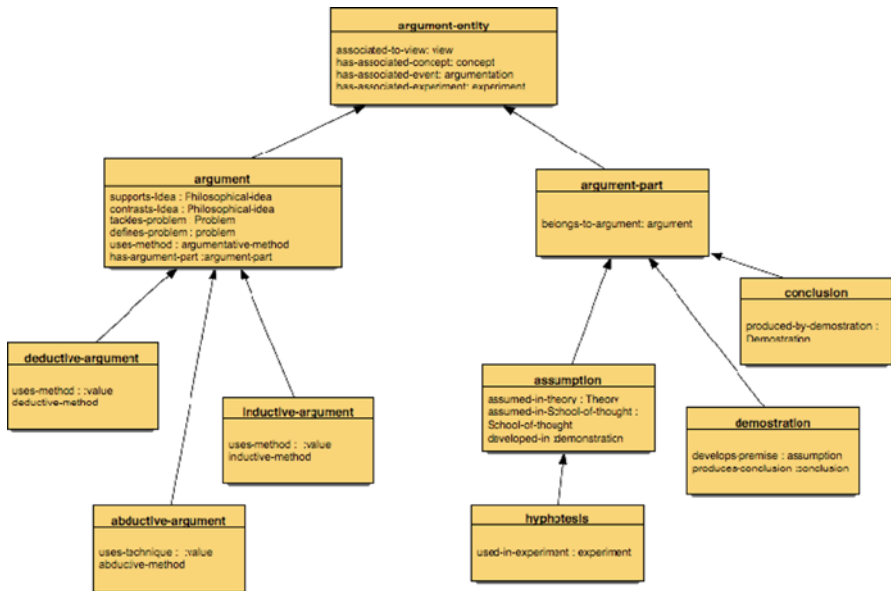


Fig. 5 Argument and argument-part

hypothesis (a subclass of *assumption* specifically referring to argumentations based on experimental evidence). It is important to note that this is only a simplified classification of the entities that can possibly build up an argument. In the future, other work from the argumentation community (Kirschner et al. 2003) could be brought in, so to represent at a finer granularity the different argument structures.

3.2 Problem-area

In order to give an account of the distinctive features of fields of study, we decided to use as a starting point a problem-centered approach. This means that we tended 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 criteria (*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 behaviour 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 the 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*.

```
(def-class Field-of-study (Problem-area)
  ((defined-by-view :type view)
   (has-exemplar-theory :type theory)
   (has-methodology :type method)))
```

As an example, we show a possible formalization of an old fashioned `field-of-study`, “phrenology”.

```
(def-instance phrenology field-of-study
  ((has-referred-author Franz-Gall)
   (defined-by-view phrenology-theory)
   (contains-problem what-is-personality-what-is-character
    relation-personality-skull)
   (has-criteria skull-shape-determines-personality-thesis)
   (sub-area-of psychology)
   (related-to-area craniometry physiognomy)
   (has-methodology phrenological-analysis)))
```

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 a practitioner for clarification), 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,

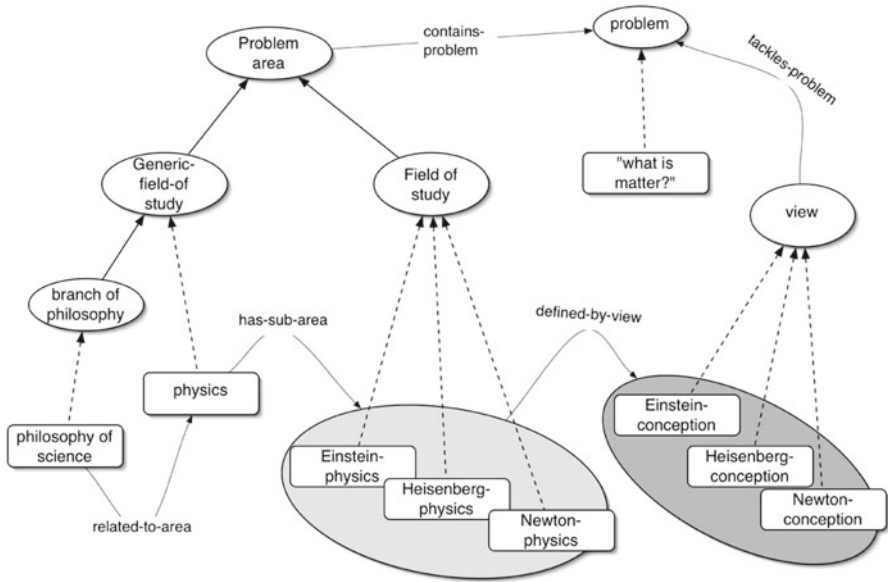


Fig. 6 Generic and specific fields of study

we also need to represent the fact that they are all part of a more generic (and probably *emptier* in its meaning) field-of-study.

Our solution to this problem (cf. Fig. 6 above) 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 we have the following OCML rule:

```

(def-rule generic-field-rule
  (defined-by-view ?GF ?V) if (generic-field-of-study ?GF)
                             (has-sub-area ?GF ?F)
                             (defined-by-view ?F ?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.

3.3 Problem

The problem class represents a very central notion in philosophy, since it is usually the point of departure of any investigation (which often culminates with the creation of a view). Examples at hand are many: we talk about the “mind-body” problem, the “alienation” problem or the “problem of the universals”. A key feature we can easily recognize is that a problem is always framed within a larger context which gives a more

precise connotation to it. So, for example, Marx considered the “alienation problem” to be rooted in “economy”, while Searle treats the “mind-body problem” within the “philosophy of mind”. Therefore, the problem exists within a problem-area. Moreover, the context which makes us understand a problem is given also by the set of assumptions that justify its existence. Or better, by the views and arguments that define it (and, conversely, try to solve it). The remaining properties of problem, as shown below, relate them to other problems or to the view and arguments that tackle them.

A special role is held by the property *has-problem-type*, which can have value “open-problem” (meaning a problem which does not have any solution), “multilemma” (a problem having or allowing multiple solutions), “dilemma” (a problem allowing two solutions only, but neither of the two being satisfactory) and “paradox” (a problem whose solutions seem equally plausible, but when considered together generate a contradiction). Essentially, these concepts describe a problem from the viewpoint of the number of solutions it has. We have modelled them as instances of the class *problem-type* (which is not in the *philosophical-idea* branch, but is instead a subclass of *CIDOC's type*), since they do not appear to be ‘essential’ for the definition of a problem, but just accidentally related to the existence of any solution. In other words, a *definitional-problem* (see below) will always maintain its structure, regardless of being an “open-problem” (i.e. having no solutions) or a “multilemma” (i.e. having various solutions).

From the analysis of the literature, we thought it was useful also to provide a classification of problems based on their *morphology*. That is, on their external structure, which can be sometimes related to their content, but is usually independent from it. In total, we identified six ‘morphological types’ of problems:

- (1) the *existence-problem* has usually the form “Does X exist?”; specializations are *existence-as-concrete-problem* (“Is X concrete/real?”) and *existence-as-abstract-problem* (“Is X abstract?”)
- (2) the *definitional-problem* has usually the form “What is X?”. Specializations are *definitional-problem-essence* (“what are the characteristic traits X has?”), *definitional-problem-attribute* (“what are the attributes X has?”) and *composition-problem* (“What is X composed of?”)
- (3) the *functional-problem* has usually the form “What is the function of X?”; the only specialization is *purpose-problem* (“What is the purpose of X?”)
- (4) the *relational-problem* has usually the form “What is the relation between X and Y?”; specializations are *dependence-problem* (“Are X and Y dependent?”), *dependence-cause-problem* (“Is X the cause of Y?”), *dependence-effect-problem* (“Is X the effect of Y?”), *independence-problem* (“Is X independent from Y?”), *equality-problem* (“Is X equal to Y?”) and *difference-problem* (“Is X different from Y?”).
- (5) the *modality-problem* is a problem about the degree of certainty X is likely to happen (or not). Specializations are *necessity-problem* (“is X necessary?”), *possibility-problem* (“is X possible?”), *contingency-problem* (“is X contingent?”) and *impossibility-problem* (“is X impossible?”)

- (6) the *factual-problem* has the form “how, in what way does X happen, or manifests itself?”.

At the time of writing, we are instantiating these problem templates by filling the empty spaces in the question with instances of concept. For example:

```
(def-instance what-is-virtue definitional-problem
  ((contains-concept virtue)
   (has-problem-type multilemma)
   (exists-in-area ethics)
   (related-to-problem what-is-value)
   (is-tackled-by-View Plato-philosophy Aristotle-
    philosophy stoic-philosophy)
   (linked-to-fact death-of-socrates)))
```

A much more interesting solution would be instead letting any instance of *philosophical-idea* be filling those spaces. This would result in a powerful reification mechanism: e.g. we could define a problem about the relation between two other problems. Moreover, we are also investigating how to use these structures for producing inferences (e.g. from a *relational-problem*, we can create a path which links to the *definitional-problems* of the concepts related). These and other issues (such as how to classify problems according to their ‘contents’ e.g. “moral problem” or “epistemological problem”) will be investigated in future research.

3.4 Method

Various ontologies introduce a class named ‘procedure’, with reference to any *sequence-like specification*. Similarly, a heuristic or method in philosophy is essentially defined as a series of steps leading from a problem towards its solution. Depending on whether the method suggests a practical activity, or an intellectual one, we classified instances as belonging to *abstract-method* or *practical-method* (see Fig. 7).

The main types of *abstract-method* are *logical-mathematical-method*, *rule-of-inference* and *argumentative-method*. The first one subsumes *algorithm* and comprises instances such as “the quick-sorting algorithm”, Wittgenstein’s “truth-table method” or Leibniz’s “infinitesimal calculus”. The second class refers to rules that are used to justify the steps in a formal proof of the validity of a more complex argument. For example, we can have “modus ponens”, “hypothetical syllogism”, “conjunction”, “double-negation elimination”, etc. The class *fallacy*, instead, refers to invalid argumentative steps that may appear convincing at first glance because they closely resemble legitimate patterns of reasoning. For example, fallacies can be the “illicit major”, “affirming the consequent”, “denying the antecedent”, “affirming the alternative”, etc. Finally, the class *argumentative-method* categorizes famous and well-established argumentation styles, such as “deductive argument”, “argumentum a fortiori”, “argumentum ad hominem”, “argumentum ad populum”, etc.

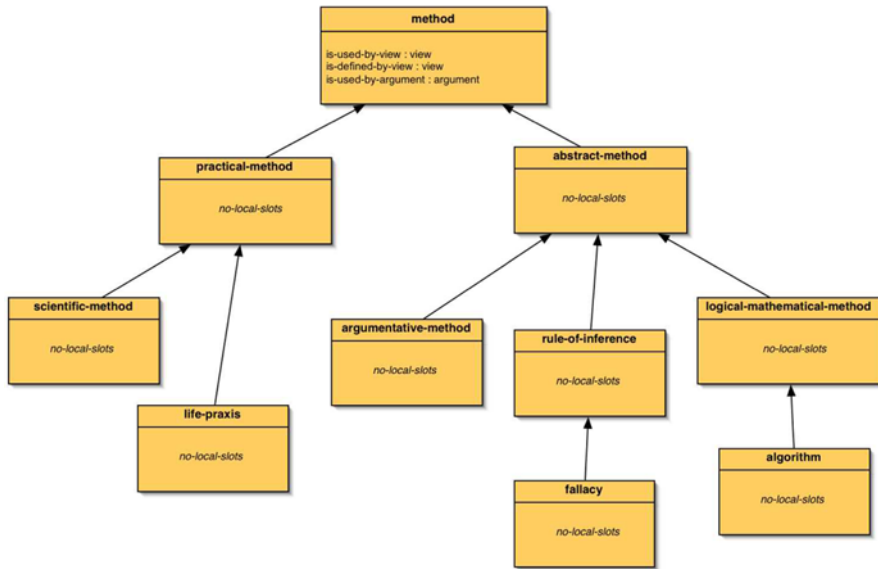


Fig. 7 The method branch of the ontology

The other branch of method, *practical-method* is divided into *scientific-method* and *life-praxis*. With the first class, we refer to any structured method to investigate reality, in a “scientific” manner (e.g. so to produce and test some explanatory hypotheses). Examples can be “Bacon’s scientific method” or “Galileo’s scientific method”. The second class instead is a method of life conduct, such as the Epicurean’s “*ataraxia*” (e.g. a description of conduct to follow in order to achieve the tranquility of the soul) or a practice of meditation in eastern philosophies.

3.5 View

This is a generic class referring to propositions 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).

Due to 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: views can *use* other ideas, *tackle* problems, *influence* and *support/contrast* each other, and *be-supported* by arguments. Most of them seemed to reflect quite well the common sense understanding of philosophy, so we will not treat them one by one here.

However, the feature we want to highlight here is how views can have varying granularities. From our analysis of the literature, we identified four possible kinds of view:

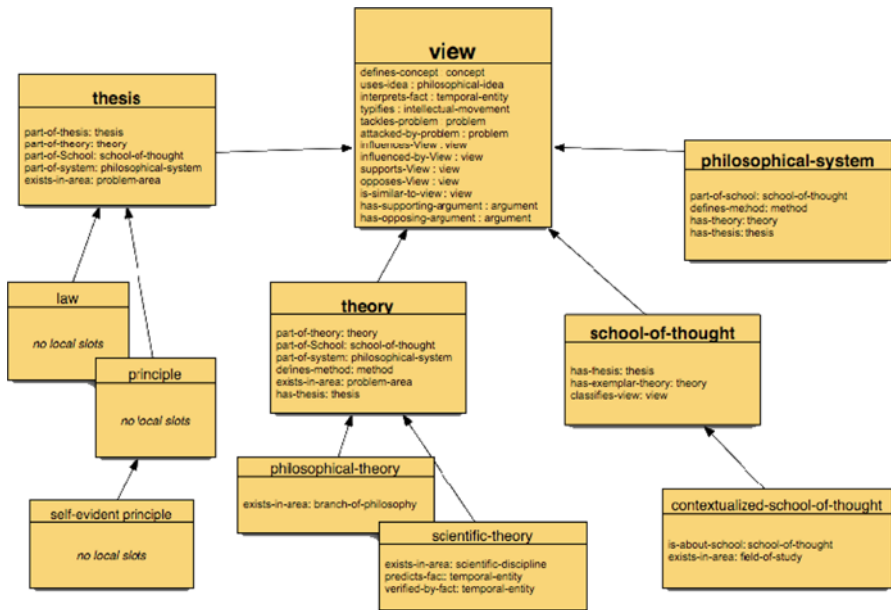


Fig. 8 The view-types

school-of-thought, theory, philosophical-system and thesis (see Fig. 8). The main differences among them depend on the degree of generality they exhibit and the level of complexity they have. In Fig. 9, we can see a small example including different views and some relations they entertain with each other. In the following four paragraphs, we will examine them one by one.

3.5.1 Thesis

A thesis 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” (as recognized by Stenius (Stenius 1960)), which can be instantiated as follows:

```

(def-instance independence-state-of-things thesis
  ((defines-concept state-of-things independence)
   (part-of-system wittgenstein-first-philosophy)
   (part-of-theory picture-theory-of-language)
   (has-string-description "State of things are
    independent of one another")))
    
```

The local properties of thesis are the *part-of* relations linking it to the other subclasses of view. Most of its properties are therefore inherited.

However, not all theses have the same status: two subclasses, law and principle, refer respectively to theses with vast predictive power, especially in scientific areas (e.g. the “law of universal gravitation”), and to theses that play a fundamental

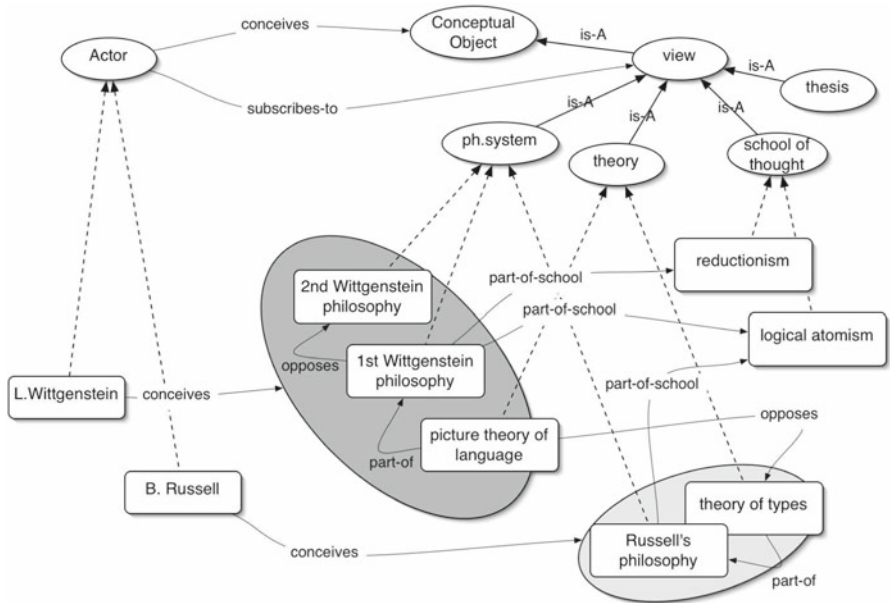


Fig. 9 The view-types instantiation

role within a view, usually a philosophical one (e.g. a principle in medical ethics). Finally, if the principle is not demonstrable but self-evident, it becomes a self-evident-principle. For example:

```
(def-instance principle-of-contradiction self-evident-principle
((defines-concept truth thought)
(part-of-system aristotle-philosophy)
(exists-in-area logic))
(has-string-description "One cannot say of something
that it is and that it is not in the same respect and
at the same time")
(appears-in Metaphysics-book-IV)))
```

3.5.2 Theory

With the class theory, we refer to a systemic conceptual construction with a coherent and organic architecture. A theory explains a specific phenomenon (or a set of phenomena) and typically answers to an already existing problem. Examples can be Darwin’s “theory of evolution” or Quine’s “verification theory”. The first one is a scientific-theory, while the second is a philosophical-theory. The main difference between the two is that the last one is not necessarily hypothetical and therefore it does not need experimental verification (although it can be provided with it).

The local slots of theory define the following properties: *part-of-theory* expresses the situation where theories are composed by other theories (e.g. Plato’s “theory of metempsychosis”, which is contained and dependent on the “theory of anamnesis”); *part-of-school* can be used to express that a theory is classified as part of a school of thought (e.g. when we say that the “picture theory of language” is a kind of “reductionism”); finally, *part-of-system* links a theory to an author’s philosophy (e.g. the “theory of eternal recurrence” is part of “Nietzsche’s philosophy”). Moreover, theories can *define-methods* (e.g. Wittgenstein’s “picture theory of language” defines the “truth tables method”), they exist within a specific problem-area (*exists-in-area*) and usually within them, we can easily identify a set of thesis (*has-thesis*).

A philosophical-theory does not differ much in its formalization from its direct super-class, apart from the fact of having range branch-of-philosophy on the property *exists-in-area*. The same property, instead, would have value scientific-area in the case of a scientific-theory. Moreover, a scientific-theory can be further defined as having a more peculiar relationship to the facts it tries to explain, as it is usually required to be verified (proved) by them, and to be able to predict them too.

3.5.3 Philosophical-system

A philosophical-system might appear as a theory, at first sight, but it differs from it essentially for its generality and breadth. 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 (which singularly taken, approach problems coming from different problem areas) which 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. So, for example, we can have “Epicurean philosophy”, “Kantian philosophy” or “Hume’s philosophy”. We must remember, however, that this class does not correspond to the mere sum of an author’s theories: in fact, thinkers might produce more than one independent system, during their lifetime (e.g. the first philosophy of Wittgenstein, as opposed to the second one).

Finally, we also recognized how a philosophical-system (although being inherently related to various problem-area) is often considered as representative of a school-of-thought (which, as explained in the next section, is instead usually related to a specific problem-area). In other words, it makes sense to say “the philosophy of Hume is scepticism”, even if, in such a case, we implicitly refer to only certain aspects of his philosophy (i.e. his epistemology). As this is a normal practice for scholars, we reckoned important also for our ontology users to be able to quickly classify philosophies using the *part-of-school* property, without having to specify the relevant theories or thesis. In order to prevent wrong generalizations (e.g. inferring that all the theories of Hume are “sceptical”), we use a set of purpose-built rules. Finally, other rules also guarantee the consistency between philosophical-systems and the theories composing them (e.g. if a theory defines a method or a concept, the philosophy comprising the theory is also considered to define them).

3.5.4 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”.

Sometimes they can be so abstract (as in the case of “monism”) that they do not imply anymore than a link to a specific problem or area, but refer only to the ‘formal features’ of the view they classify. For example, in the case of “monism”, what is implied is just ‘a view that admits only one principle as fundamental’.

A school-of-thought, compared to the other views, is not as formalized and specific as a theory, and not as broad and systematic as a philosophical-system. Accordingly, in our model, we decided to limit its contents to instances of thesis. Because of this “generic” flavour, we often perceive the meaning of schools as being vague and abstract (e.g. when trying to specify what is “rationalism”).

On the contrary, we noticed that this is not the case when we refer to (1) their “instantiation” within a problem area (e.g. the “ethical rationalism”) and (2) their specific “expression” within an author’s philosophy (e.g. the “rationalism of Kant”). These last two examples seemed to us quite important; therefore, we attempted to give an account of them also in the ontology.

According to our analysis, the first case (“ethical rationalism”) relates to the fact that schools of thought often have a ‘contextualized’ version. That is, they assume a different and more specific meaning when associated to a specific problem-area. For example, “rationalism”, can be found in “epistemology”, in “ethics”, in “meta-physics” or in “philosophy of religion”. The interesting phenomenon, in this case, is that the contextualized versions do not always have much in common and sometimes are even surprisingly unrelated. For example, let us mention the different meanings of “cognitivism” in “psychology” and in “meta-ethics”. Therefore, in order to keep separated the meaning of generic schools of thought from their localized ones, we introduced the class contextualized-school-of-thought, which has the additional slot *exists-in-area* with range *field-of-study*.

Instead, regarding the second case (the “rationalism of Kant”), we concluded that it refers to the fact that schools of thought are normally used as ‘classifiers’ of other views. We showed in a precedent paragraph how this relation is already captured by the *part-of-school* property of theory and *philosophical-system*. In a similar fashion, we created also the slot *has-exemplar-theory*, which refers to the theory that inspired the school-of-thought, and is likely to help in understanding its original sense.

3.6 Rhetorical figure

With this class, we aimed at grouping figures of speech or statements embodying some rhetorical value; usually these objects of discourse are used for emphasis, for clarity or as a device in the philosophical argumentation. Many of these entities could also

have fitted as subtypes of *argument-part*, since in most cases they play that role. However, since often they assume a singular significance in the history of thought (i.e. the “myth of the cave”), we decided to represent them separately, so that they could be treated (and re-used) as independent entities.

We have defined three types of *rhetorical-figure*: *metaphor*, which subsumes *myth* and *analogy*; *maxim-motto*, and *thought-experiment*. All of them can be described by using the properties *used-in-argument* and *used-in-view*.

Examples of the first type is the aforementioned “myth of the cave”, or Hegel’s metaphor of the “night, in which all cows are black” (used in the argument against Schelling). *Maxim-motto* refers instead to famous and exemplar statements or expressions philosophers used to sum up their position. For example, Descartes’ “*cogito ergo sum*”, Hobbes’ “*homo homini lupus*” or the ancient maxim “*ex nihilo nihil fit*”. Finally, *thought-experiment* refers to mind-simulations used to prove a point: among them, we can remember Searle’s Chinese-room thought-experiment (used to attack strong AI), Putnam’s twin-earth thought-experiment (used to support “semantic externalism”) or David Chalmers’ “unconscious zombies” thought-experiment (used to attack “physicalism”).

3.7 Concept

A concept is an atomic element (i.e. not further decomposed) in the ontology. Instances of concepts can be “ego”, “evolution” or “god”. In determining what is a concept, we are not interested in its cognitive and linguistic features (i.e. the fact that it carries one propositional content, or that it is expressible through one or two words), but mostly in its functional role within the economy of a philosophy or a theory. That is, we tend to see a concept as an element which is defined by a view as primitive, and which is in a net of relations with other concepts.

According to a ‘philosophy of minimum commitment’, we have chosen not to formalize specific philosophical concepts as classes, but to provide means to create alternative interrelated nets of instances which could resemble (and could be exported as) a small taxonomy. Thus, the creation of a network of interrelated concepts relies totally on the annotator. We expect people to organize the knowledge associated with an author’s conception very differently, according to user needs, background and interests.

A concept can be linked to other concepts through various relations: specialization and generalization (*is-specialization-of* and *is-generalization-of* properties); similarity of meaning (*is-equivalent-to*), e.g. for the concepts “inexpressible” and “ineffable” in Wittgenstein; antinomic contrast (*has-opposite-concept*), e.g. when two concepts are part of a dichotomy; generic semantic closeness (*has-related-concept*), e.g. when they concur in explaining the same phenomena; notional dependency (*requires-concept*), e.g. with concepts such as “buy” and “pay”; causation (*causes-concept*), e.g. with concepts such as “to kill” and “to die”.

For example, the Wittgensteinian concept of “picture” could be defined as follows:

```
(def-instance picture-by-first-wittgenstein concept
  ((has-common-name picture)
   (defined-by-view first-wittgenstein-philosophy)
   (is-specialization-of fact-by-first-wittgenstein)
   (is-generalization-of logical-picture-by-first-wittgenstein)
   (has-similar-meaning-as picture-by-hertz)
   (is-in-contrast-with)
   (is-in-relation-with isomorphism-by-first-wittgenstein
    form-of-representation-by-first-wittgenstein represent-
    ing-relation-by-first-wittgenstein))))
```

Finally, the *has-common-name* property (whose range is *idea-appellation*) is used for separating the *concept* object from the name used to identify it (e.g. “picture” in English, “immagine” in Italian, “image” in French, etc.). Let us remind that CIDOC provides a useful facility to detach entities from their names, that is the *appellation* class (it is located in the *persistent-item* branch of the ontology). By instantiating this class, for example, we can define multiple names for the same place or for the same person. Analogously, we added also an *idea-appellation* class in order to support the separation of an *idea-object* from its names.

This turned out to be quite a handy feature, because often there are no explicit properties stating the relationships between two instances of *concept*, but the fact that they have the same name. In Fig. 10, we can see an example of how the word “alienation” (which is an *idea-appellation* instance) could be referring to four different concepts. Each of them, in fact, is defined by a different *view*, categorized by different *school-of-thought* and typical of different *problem-areas*.

3.8 Distinction

We have a *distinction* when two ideas or more stand out as particularly meaningful in their opposition. That is, the specificity of their sense is obtained or clarified by their being different, but complementary. For example, “Hume’s distinction between truth of reason and matters of fact”, “Aristotle’s distinction between essence and accident”, or “Frege’s distinction between extension and intension”. Together, the two concepts fill up a whole, with respect to a specific domain of reference, e.g. “epistemological” (regarding the limits of human knowledge) or “ontological” (regarding the structure of being). A distinction can have an arbitrary number of concepts (e.g. “Aristotle’s four types of causes”), but when comprising two concepts only, it is also called *dichotomy*. For example:

```
(def-instance hume-fork dichotomy
  ((has-referred-author david-hume)
   (related-to-area epistemology)
   (related-to-problem what-can-we-know)
   (defined-by-view hume-philosophy)
   (contains-concept relation-of-ideas matter-of-fact))))
```

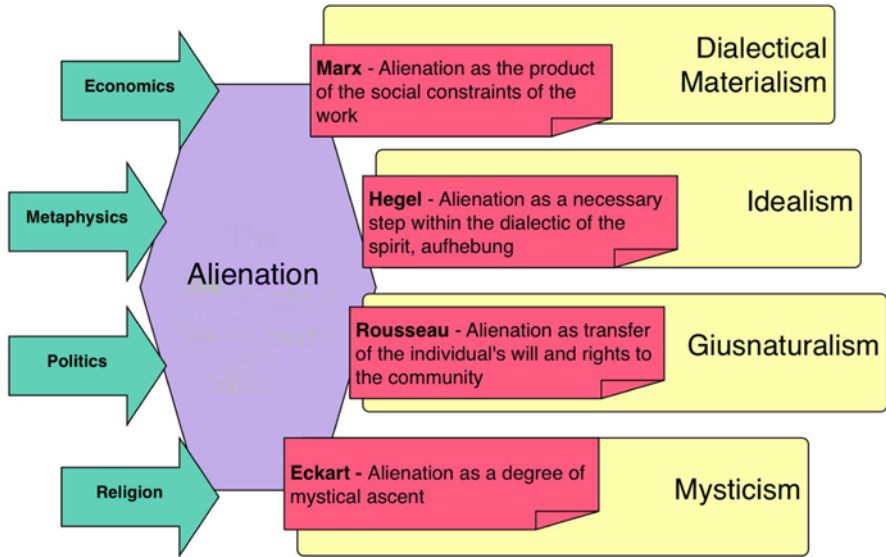


Fig. 10 The four concepts behind the philosophical term “alienation”

4 Putting things together: the PhiloSurfical tool

In this section, we describe the main features of PhiloSurfical⁶, a prototype software that allows the navigation of a semantically enhanced version of Wittgenstein’s *Tractatus Logico-Philosophicus* (Wittgenstein 1922). By relying on the various levels of abstraction provided by the ontology, the software lets users browse the text and other associated resources in a *contextual* manner. For example, users can select all text instances which have been annotated with a specific concept, discover how this concept relates to other concepts in the philosophy of Wittgenstein and, in general, access data using the network of relations that have been formalized in the ontology (Fig. 11).

This methodology, which has been previously defined as *ontology-based navigation* (Crampes and Ranwez 2000), can be further developed by means of an approach modelled on *narratology* (Chatman 1978). As already discussed in an earlier publication (Pasin and Motta 2005), 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 one based on the most relevant *schools of thought*. Similarly, the same set of philosophical ideas

⁶ The application is available online at <http://philosurfical.open.ac.uk>

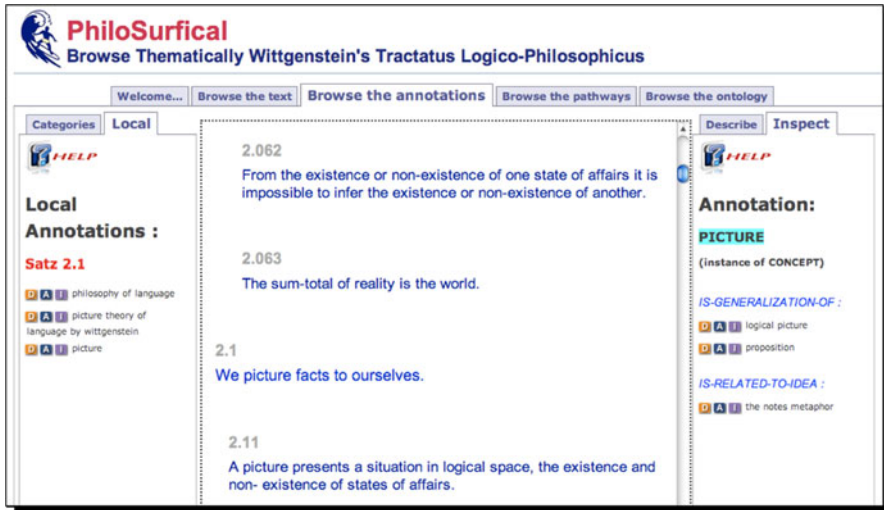


Fig. 11 Screenshot of the PhiloSurgical application

could be organized differently if investigated under a *problem-centered* perspective, a *theory-centered* one, or simply one based on their *historical* succession.

In other words, our approach takes the notion of a ‘digital narrative’⁷ (Brooks 1997) and attempts to transpose it to the specific scenario made up of philosophical entities. Accordingly, with PhiloSurgical, we aimed at creating a virtual environment for exploring user-triggered digital narratives, which we also call *learning pathways*.

Due to space limitations, we cannot give here a complete description of all the pathways made available in PhiloSurgical. In order to better understand the role and usage of the ontology within the software tool, we will instead focus on the construction of a Tractatus-related knowledge base and on the functioning of a specific type of learning pathways, the *theoretical* ones (cf. Table 1 below).

4.1 Creating a knowledge base for the Tractatus

Although the ontology was created with the aim of facilitating 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), PhiloSurgical strongly relies on an internal knowledge base of our creation.

Before going further, an important clarification has to be made. By instantiating the ontology with various Tractatus-related data, we inevitably created a ‘unified’ philosophical view of this text, in the sense that we had to privilege certain interpretations instead of others. Certainly, such a result is not representative of the reality,

⁷ Brooks defines it as a “system of specially stored and organized narrative elements which the computer retrieves and assembles according to some expressed form of narration”.

Table 1 The theoretical pathways available in PhiloSurgical

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. For example, 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 which are all specializations of the same generic one. This pathway is related to the formalization presented in Sect. 3.5.4: e.g., by focusing on ‘atomism’ we would be able to see the related <i>contextual</i> 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. For example, starting from ‘Wittgenstein’s theory of language’, we could go to ‘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 <i>interpretations</i> 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.

where the amount of critical literature on this influential text is just enormous. Thus, consistently with what was emphasized in Sect. 1, our views on the *Tractatus* have no pretension whatsoever to be representative of all the literature, or to be truer than others. In general, we just aimed at creating a pedagogical resource that could be used as an *introduction* to 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 our ‘learning pathways’ approach.

It is also useful to point out that in this respect, our work differs radically from other digital editions of Wittgenstein’s works, e.g. Bazzocchi’s *Tractatus* (Bazzocchi 2007) or the famous Bergen edition of the *Nachlass* (Pichler 2002). Our aim was simply to *test the quality of the ontology* by instantiating it with real-world philosophical data. The other digital editions focus instead on creating a *new version* of a classic text, usually by taking advantage of various features of the digital medium.

The key difference here is that our research interest concerns the modeling and integration of philosophical data in an open context like the Semantic Web. Within such a scenario, the *Tractatus* is for us just a ‘handy’ testbed for the instantiation

of the ontology (first of all, because it is a highly structured text, thus simplifying the analytical task of dissecting it into meaningful units). On the contrary, the digital editions mentioned above do not make available the (implicit) semantic model used in building the application, that is, they do not present it in the form of an ontology that others can reuse, modify or employ for exchanging philosophical data.

The creation of PhiloSurgical's knowledge base is composed of three phases.

- (1) The transformation of the text itself into a semantic format. First, we downloaded the Gutenberg edition of the Tractatus⁸, which corresponds to the English translation made by David Pears and Brian McGuinness in 1961. We then built a suitable parser to extract the different paragraph numbers and text, so to populate the relevant parts of the ontology (mainly, subclasses of `information-objects` representing text entities at various levels of abstraction). Moreover, we repeated this process with two other editions of the text, the translations made in 1922 by Charles Kay Ogden⁹ and the original German version¹⁰. As a result, we created 1591 instances representing Tractatus sentences.
- (2) The annotation of the text's paragraphs. For the annotation phase, we worked in collaboration with a Wittgenstein scholar, Andrea Bernardi. Essentially, we went through all of the text's paragraphs with the purpose of extracting the key-concepts they are dealing with. We then drew a map where it is possible to see the association of each concept to the paragraphs where it is mentioned. During this process, our philosophy expert also created some basic relations that contextualize the concepts with respect to one another, so to form links among them (*inclusion, opposition, similarity...*). Moreover, we annotated a number of specific relationships the concepts entertain with other types of philosophical entities (e.g. a theory *belongs to a school of thought*, a theory *defines a concept*, an author *belongs to a philosophical school*, etc.). To conclude this process, we generated a layer of interpretation instances about the Tractatus (analogous to what described in Sect. 2.3). By using this method, we created a total of 639 instances representing interpretations of Tractatus sentences, 434 instances of philosophical ideas related to the text and 290 interpretations of the ideas.
- (3) The enlargement of the knowledge base through the addition of further philosophy-related instances. This was done automatically, mostly by 'scraping' the relevant information from websites in the public domain. Afterward, this data was evaluated and sometimes refined manually. In general, we imported data about famous philosophers (more than 7,000 instances of person), schools of thoughts (about 500 instances of `school-of-thought`), the secondary Wittgensteinian literature (about 100 instances of `information-object`) and philosophical dictionary entries (about 5,000 instances of `information-object`).

⁸ <http://www.gutenberg.org/etext/5740>

⁹ <http://www.kfs.org/~jonathan/witt/tlph.html>

¹⁰ <http://www.tractatus.hochholzer.info/>

The screenshot displays a web-based pathway system. On the left, the 'Pathways' section includes a 'Recent Items & Search' tab and a 'HELP' icon. Below this, the 'Item in focus' box contains the text 'FOUNDATIONS-OF-MATHEMATICS-PROBLEM' with a 'change | instance info' link. The 'Pathways list' section features a scrollable list of pathway options, with 'Problem-centric map of the attempt to solve a problem' highlighted in blue. To the right of this list, a 'Type of pathway:' section indicates 'THEORETICAL' and provides a description: 'This pathway takes a problem instance and retrieves information related to the competing views (theories, schools of thought, philosophies) that tackle and attempt to solve the problem'. The 'Results' panel on the right, titled 'Results' with a '[see in a graph]' link, shows a list of related items, including 'IS-TACKLED-BY-VIEW', 'PLATONIST MATHEMATICAL REALISM', 'HAS-MAIN-EXPONENT', 'HAS-EXEMPLAR-THEORY', 'OPPOSES-VIEW', and 'MATH-INTUITIONISM-MATH-INT'. Each item is accompanied by a brief description of its philosophical context.

Fig. 12 Pathway representing the various attempts to solve a problem

4.2 Ontology-enabled pathways for learning

A ‘pathway’ is essentially a way to retrieve different instances stored in the knowledge base and organize them into a coherent whole. We classified pathways according to the ontological type of their ‘entry point’ (i.e. the instance we start the pathway from), 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.

From the point of view of a learner, such mechanisms can be used as follows. First of all, users select a content of interest as the starting point of a pathway (Fig. 12, ‘item in focus’ box). Learners may then click on one of the available choices appearing in the ‘pathways list’ panel (see Fig. 12, 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. 12, ‘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.).

For example, starting from the *problem* instance called “problem of the foundations of mathematics” we might select the ‘problem-centric map of the attempts to solve a problem’ pathway. As shown in Fig. 12, 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. *has-main-exponent*, *has-exemplar-theory*, etc.).

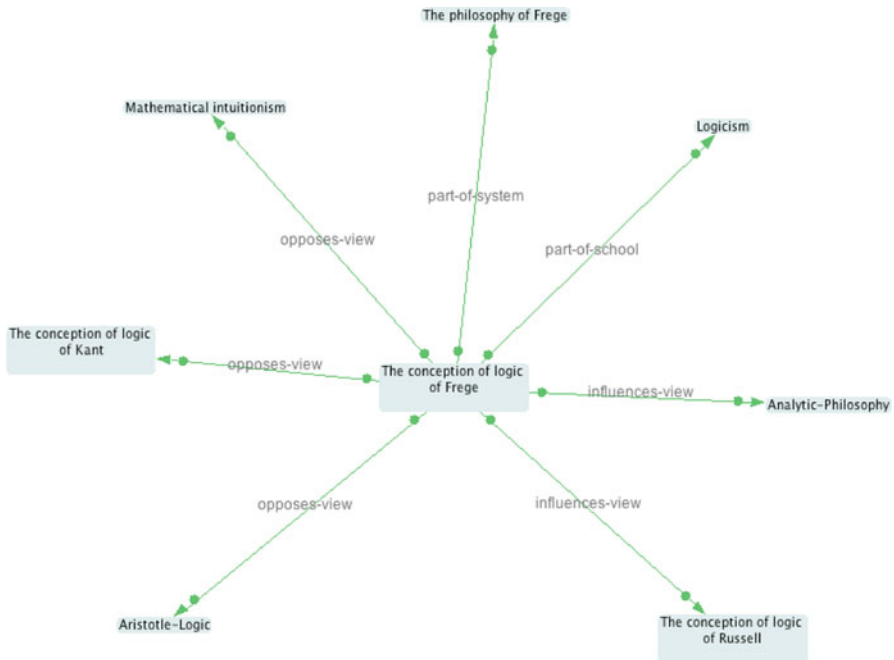


Fig. 13 Graphical view of a theoretical pathway about Frege

Furthermore, by clicking on the ‘see in a graph’ button learners can view the pathway’s results using a graphical visualization. For example, in Fig. 13, 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.

Internally, PhiloSurfical represents pathways as abstract procedures applicable to any ontology-compliant data repository. For instance, in Fig. 14, 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 1 above).

In general, after a pathway is triggered, we scan the knowledge base for instances of *interpretation* mentioning the item which has been selected by the user. Subsequently, we analyze the *interpretation* instances retrieved for the purpose of finding information which is relevant to the specific pathway the user has selected. For example, in the case of ‘influences among related views’, we are interested in relations such as *supports-view* and *opposes-view*. If some results are found, we store them for the visualization phase. Of course, each pathway presents individual differences too: e.g. the ‘problem-centric map of the attempts to solve a problem’ pathway searches for relevant *interpretation* instances twice: first with a *problem* instance, second with a *view* instance; instead, ‘influences among related views’ is a function that calls itself *recursively* a predefined number of times, so to create a ‘nested’ map of related views.

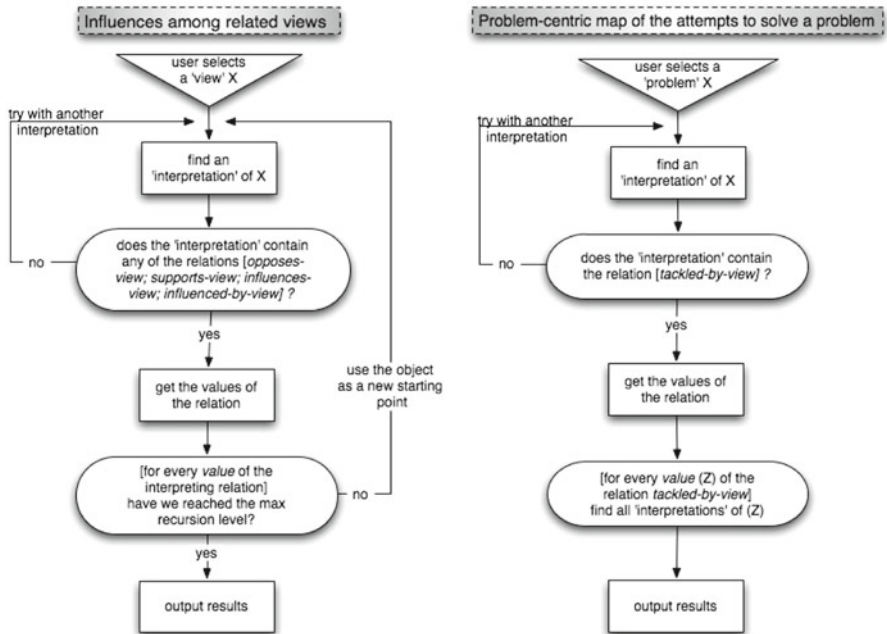


Fig. 14 Abstract representation of two pathways' algorithm

Usually, the output of these algorithms is a very 'concise' representation of the final dataset we present to the user. For example, the results of a theoretical pathway involving different concepts related to the same author may omit the repetition of the *is-author-of* property. On the contrary, the data shown in the user interface need to explicitly mention all these relations.

At the moment, this 'explosion' process is handled by two routines, depending on whether the results are presented in *html* or in the *java-based* graphical view. In future releases of PhiloSurgical, it is likely that we will add also other types of data visualizations.

5 Related work

The most relevant (and to our knowledge unique) attempt to systematically formalize the philosophical domain is the one carried out in Niepert et al. (2007), as part of a larger 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 (Gábor et al. 2005; Vieira and Ciula 2007). In this context, the *Discovery project* (2008) stands out for its explicit goal of creating a Semantic Web infrastructure specifically for *philosophers*. From the ontological point of view, the authors plan to use a ‘network’ of ontologies (Nucci et al. 2007). 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 ideas (and especially philosophical ideas), we found no evidence of relevant work in the *knowledge representation* research literature. Although models such as Wordnet (Fellbaum 1998) and Cyc (Lenat and Guha 1990) 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. The noteworthy exception here is the DnS module of Dolce (Gangemi and Mika 2003), 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). Furthermore, our formalization of fields of studies (cf. Sect. 3.2) 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 (Welty and Jenkins 1999) could be well suited also for the philosophical domain. We plan to investigate further this issue in future work.

Finally, it is worth mentioning recent research aimed at facilitating the semantic navigation of digital resource 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* (Ranganathan 1990), have been tested in various humanities domains, such as classical music (Schraefel 2005), visual arts (Hildebrand et al. 2006), cultural heritage (Hyvonen et al. 2008) and literature (Nowvieskie 2005). 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.

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 provided feedback and support during the

various stages of the research. In particular, Andrea Bernardi, Keith Frankish, Gordon Rugg, Marian Petre, Riichiro Mizoguchi and Martin Doerr.

References

- AKT. Reference ontology v.2—AKTive portal ontology v.2. (2002). <http://d3e.open.ac.uk/akt/2002/portal-ocml-v2.0/portal-ocml-v2.0-t.html>.
- Allen, J. F. (1984). Towards a general theory of action and time. *Artificial Intelligence*, 23, 123–154.
- Bachelard, G. (1938). *La formation de l'esprit scientifique*.
- Bazzocchi, L. (2007). On butterfly's feelers: Some examples of surfing on Wittgenstein's Tractatus. *Proceedings of the 30th International Ludwig Wittgenstein-Symposium*.
- Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The semantic web. *Scientific American*.
- Brooks, K. M. (1997). Do story agents use rocking chairs? The theory and implementation of one model for computational narrative. *Fourth ACM multimedia conference*.
- Chatman, S. (1978). *Story and discourse*. Ithaca: Cornell University Press.
- Crampes, M., & Ranwez, S. (2000). Ontology-supported and ontology-driven conceptual navigation on the world wide web. *11th ACM hypertext conference*.
- Crofts, N., Doerr, M., Gill, T., Stead, S. & Stiff, M. (2005). CIDOC-CRM Version 4.2—Reference Document.
- Discovery project, official website. (2008). <http://www.discovery-project.eu/>.
- Doerr, M. (2003). The CIDOC conceptual reference module: An ontological approach to semantic interoperability of metadata. *AI Magazine Archive*, 24, 75–92.
- Farquhar, A., Fikes, R., & Rice, J. (1996). The ontolingua server: A tool for collaborative ontology construction. Stanford Knowledge Systems Laboratory Technical Report
- Fellbaum, C. (Ed.). (1998). *WordNet: An electronic lexical database*. Cambridge, MA: MIT Press.
- Gábor, Nagypál, Richard, Deswart, & Oosthoek, J. (2005). Applying the semantic web: The VICODI experience in creating visual contextualization. *Literary and Linguistics Computing*, 20, 327–349.
- Gangemi, A., & Mika, P. (2003). Understanding the semantic web through descriptions and situations. *International conference on ontologies, databases and applications of semantics (ODBASE)*.
- Gangemi, A., Guarino, N., Masolo, C., Oltramari, A., & Schneider, L. (2002). Sweetening ontologies with DOLCE. *13th international conference on knowledge engineering and knowledge management (EKAW02)*.
- Gruber, T. (1993). In Guarino, N. & Poli, R. (Ed). *Formal ontology in conceptual analysis and knowledge representation*. Boston: Kulwer Academic Publishers.
- Hildebrand, M., van Ossenbruggen, J., & Hardman, L. (2006). /facet: A Browser for Heterogeneous Semantic Web Repositories. International Semantic Web Conference—ISWC2006.
- Hyvonen, E., Ruotsalo, T., Haggstrom, T., Salminen, M., Junnila, M., Virkkila, M., Haaramo, M., Kauppinen, T., Makela, E., & Viljanen, K. (2008). CultureSampo—Finnish culture on the semantic web. The vision and first results. In K. Robering (Ed.), *Information technology for the virtual museum*. Berlin: LIT Verlag.
- Kalfoglou, Y., & Schorlemmer, M. (2003). Ontology mapping: The state of the art. *The Knowledge Engineering Review*, 18, 1–31.
- Kirschner, P., Shum, S.B., & Carr, C. (2003). *Visualizing argumentation: Software tools for collaborative and educational sense-making*. London: Springer-Verlag.
- Lenat, D. B., & Guha, R.V. (1990). *Building large knowledge-based systems: Representation and inference in the Cyc project*. Boston, MA: Addison-Wesley.
- Mizoguchi, R. (2004). Tutorial on ontological engineering—part 3: Advanced course of ontological engineering. *New Generation Computing*, 22, 198–220.
- Motta, E. (1999). *Reusable components for knowledge modelling—principles and case studies in parametric design problem solving*. The Netherlands: IOS Press.
- Mulholland, P., Collins, T., & Zdrahal, Z. (2004). Story fountain: Intelligent support for story research and exploration. *9th international conference on intelligent user interface*.
- Niepert, M., Buckner, C., & Allen, C. (2007). A dynamic ontology for a dynamic reference work. *Joint conference on digital libraries (JDCL-07)*.
- Nowviskie, B. (2005). COLLEX: Semantic collections & exhibits for the remixable web. <http://www.nines.org/about/Nowviskie-Collex.pdf>.

- Noy, N. F., & McGuinness, D. L. (2001). *Ontology development 101: A guide to creating your first ontology*. Stanford Knowledge Systems Laboratory Technical Report.
- Nucci, M., David, S., Hahn, D., & Barbera, M. (2007). Talia: A framework for philosophy scholars. *SWAP 2007, the 4th Italian semantic web workshop*.
- Pasin, M., & Motta, E. (2005). Semantic learning narratives. *International workshop on applications of semantic web technologies for e-learning (SWEL)*.
- Pasin, M., & Motta, E. (2007). Supporting philosophers' work through the semantic web: Ontological issues. *Fifth international workshop on ontologies and semantic web for e-learning (SWEL-07)*.
- Pasin, M., Motta, E. & Zdrahal, Z. (2007). Capturing knowledge about philosophy. *International conference on knowledge capture (KCAP'07)*.
- Pichler, A. (2002). Encoding Wittgenstein. Some remarks on Wittgenstein's Nachlass the Bergen Electronic Edition, and future electronic publishing and networking. *trans. Internet-Zeitschrift für Kulturwissenschaften 10*.
- Ranganathan, S. R. (1990). *Elements of library classification*. : South Asia Books.
- Rugg, G., & McGeorge, P. (2005). The sorting techniques: A tutorial paper on card sorts, picture sorts and item sorts. *Expert Systems*, 22, 94.
- Schraefel, M. C. et al. (2005). The mSpace classical music explorer: Improving access to classical music for real people. *V MusicNetwork Open Workshop: Integration of music in multimedia applications*.
- Stenius, E. (1960). *Wittgenstein's "Tractatus": A critical exposition of the main lines of thought*. Oxford: Blackwell Publishers.
- Vieira, J. M., & Ciula, A. (2007). Implementing an RDF/OWL ontology on Henry the III fine rolls. *OWLED (ESWC 07)*.
- Vygotsky, L. (1978). *Mind in society: Development of higher psychological processes*. Cambridge: Harvard University Press.
- W3C. OWL web ontology language overview. (2004). <http://www.w3.org/TR/owl-features/>.
- Welty, C., & Jenkins, J. (1999). An ontology for subject. *Journal of Data and Knowledge Engineering*, 31, 155–181.
- Wittgenstein, L. (1922). *Tractatus logico-philosophicus*. London: Routledge & Kegan Paul.
- Zúñiga, G. L. (2001). Ontology: Its transformation from philosophy to information systems. *Formal Ontology in Information Systems (FOIS)*.