2. Learning and navigation through semantic technologies: the state of the art

2.1 Introduction

In this chapter we review the literature relevant to our research. The scope of this review is quite vast, as it spans from the theoretical analysis of pedagogical theories to the technical description of Semantic Web systems used in educational scenarios. Nonetheless, broadly speaking, the chapter can be divided into two main sections.

a) In the first one we provide the reader with the theoretical building blocks necessary for understanding the *pedagogical* assumptions guiding the *design* of our system. More precisely, this section is organized as follows:
    1. Section 2.2.1 discusses the major features of constructivist learning theories.
    2. Section 2.2.2 analyzes the specific characteristics associated with learning philosophy.
    3. Section 2.2.3 introduces the theory of learning through story-construction.
    4. Section 2.2.4 outlines the basic principles of narratology.
    5. Finally, section 2.2.5 discusses the theory and design of digital narrative systems.

b) Having built up the necessary theoretical background, in the second part of this review we look at the emerging Semantic Web technologies and give an account of the major innovations they bring to the field of e-learning. In particular, we focus on systems that enhance learners' experience by means of advanced navigation and narrative-creation mechanisms. The analysis is organized in the following sections:

   1. Section 2.3.1 gives an overview of the Semantic Web (SW) vision and technologies.
   2. Section 2.3.2 summarizes a number of frameworks regarding the application of SW technologies in e-learning (SWEL).
   3. Section 2.3.3 gives various examples of SWEL systems.
   4. Section 2.3.4 provides examples of SWEL tools for the authoring and management of educational contents.
   5. Section 2.4.1 introduces the concept of 'semantic navigation' and discusses a number of systems implementing it.
   6. Section 2.4.2 discusses a recent variant of semantic navigation called 'faceted browsing'.
   7. Section 2.4.3 analyses systems employing narratology-inspired concepts for the automatic creation of hypermedia presentations.
   8. Section 2.4.4 reports on systems enhancing documents' navigation by means of semantic hyperlinking.

Finally, in section 2.5 we draw some conclusions and outline the research gap emerging from this literature review.

Figure 2-1 shows a graphical representation of this chapter's narrative structure. In particular, the three layers at the bottom (*theoretical building blocks*, *SW technologies*, *SW technologies in e-learning and navigation*) have been organized both to highlight and to provide the necessary background for the definition of the research gap.

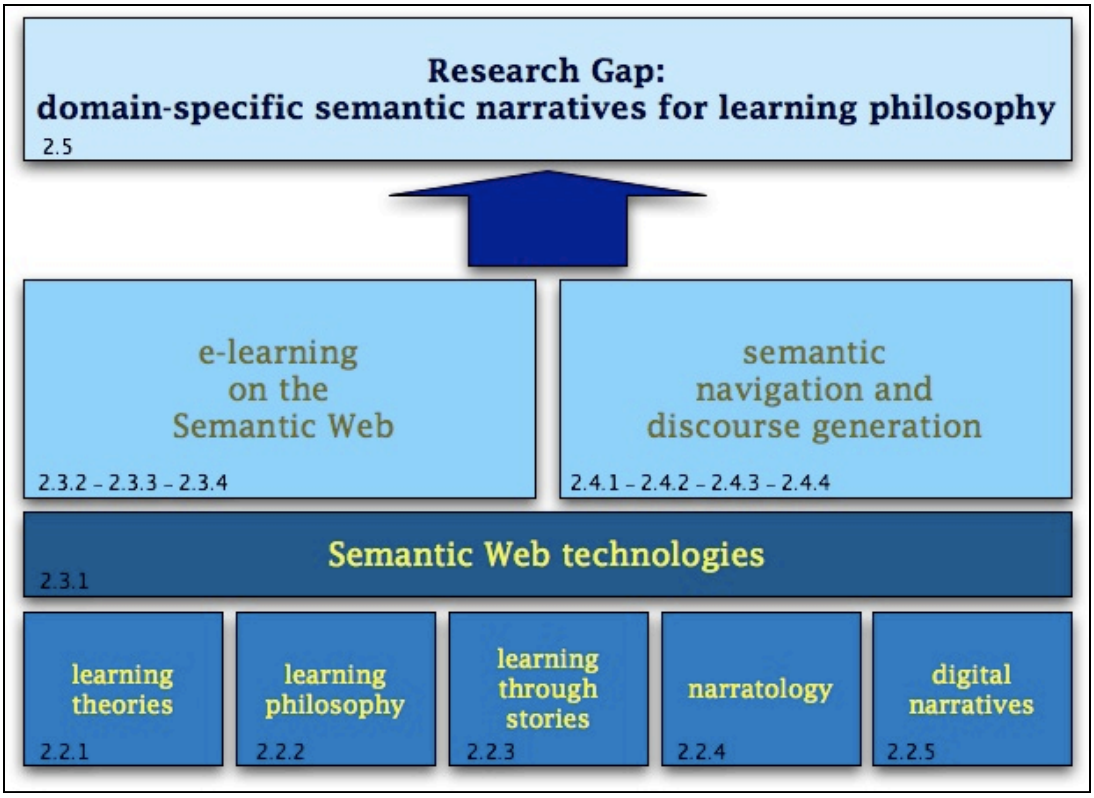

Figure 2-1 - Conceptual schema of the topics discussed in the literature review (the numbers refer to the relevant sections in this chapter)

2.2 Building blocks: learning and narratives

This section presents the main tenets of the pedagogical theories inspiring our work. In particular, we will focus on the theories describing the learning of 'abstract concepts' (such as the common concepts in philosophy) and on the connections between learning and storytelling activities. Moreover, we will give an account of the main theses underlying the field of narratology, and describe how such ideas have been transposed to the computer medium leading to the design of *digital narrative systems*.

2.2.1 Theory of learning

The learning sciences are a wide field, drawing influences from an even wider set of disciplines and from the whole history of philosophical thought (Noddings, 1998). In this section we cannot give an exhaustive account of all the standpoints and the possible approaches, but we just revisit some recent conceptions that have influenced the modern educational scenario, very much supported and at the same time determined by the use of computers.

In particular, we will focus on the theories of *cognitive apprenticeship* and *constructivism*.

The first learning theory we want to mention was firstly formulated by Brown, Collins and Duguid (Brown et al., 1989) and it is often dubbed as ‘epistemology of practice’ or ‘theory of cognitive apprenticeship’. These researchers have done an inspiring work that compares the learning activity during normal apprenticeship and the one that happens within a scholastic environment.

Usually, this difference is reduced to the difference between the categories "know how" and "know that". Nevertheless, they argue that this dualistic position is artificial, and that it should be overcome by a new epistemological standpoint. In fact, the activity in which knowledge is developed cannot be separated from its outcomes. As a consequence, knowledge cannot be treated as an abstract entity and it cannot be transferred as it was a material good.

Of course, as stated in the following passage, this dualistic epistemology has been deeply influencing also the educational practices:

> "For centuries, the epistemology that has guided educational practice has concentrated primarily on conceptual representation [...] An epistemology that begins with activity and perception, which are the first and foremost embedded in the world, may simply bypass the classical problem of reference - of mediating conceptual representations."

So, on the side of the traditional epistemology of '*possession*', limited to the exploration of the "know that", a new epistemology of '*practice*' (Cook and Brown, 1999) should develop and re-integrate the role of world and action in the knowledge generation process.

The notion of *context* is crucial in order to understand this position. We can summarize it this way: since the social and physical environment always influences the way knowledge is produced, in order to positively transfer knowledge, the same environment must be taken into account and used as a grounding for the learning process. As Brown and colleagues say, *knowledge is always situated*.

Many are the examples and case studies the authors provide in order to support this position. For instance, they did some experiments on vocabulary teaching in which students, although given very precise definition of new words, would still have great difficulties in using them within normal conversations. The knowledge acquired, explain the authors, remains abstract and does not have any link to the possible context of usage.

Following this line, the conclusion reached by the authors is that *conceptual* knowledge is to be treated as a *tool we use in everyday activity*, and that the usage of these tools is encoded in the culture of a community. Effective learning is apprehending how to use these tools meaningfully, thus *effective learning is basically a process of acculturation and experience sharing* within a group of people. The process in which conceptual knowledge is successfully used and put into practice is called by the authors 'authentic activity', since it constitutes the only setting capable of generating a real learning experience.

This epistemological viewpoint has many implications for that regards the conception of 'traditional' educational scenarios. In fact, the critique towards contemporary schooling is that it can be abstract, detached from real problems and unaware of the necessity of the students to engage with the relevant domain culture that has actually generated what is being taught. For example, claim the authors, the teaching of *mathematics* as a formalistic set of rules and methods usable only within abstract scenarios leaves the subject unlinked to the real problems that started the research. Moreover, it encourages a culture of math phobia rather than one of authentic math activity. Thus, just like a generic practitioner during a practical work activity learns 'socially' how to use a tool, analogously the *environment* has to be central in the teaching of conceptual knowledge, has to ground the intellectual work, and has to provide the application scenario of the knowledge generated.

Cognitive apprenticeship is the new category that describes a learning method based on social interaction and activity, in way similar to that evident in craft apprenticeship. For example, math teaching should be contextualized and linked to the problem solving activities that generates the abstract definitions and rules. Similarly, philosophy teaching must be linked to the real word facts where the initial questioning begins.

The benefits of this approach are well described by the authors:

> “By beginning with a task embedded in a familiar activity, it shows the students the legitimacy of their implicit knowledge and its availability as scaffolding in apparently unfamiliar tasks. By pointing to different decompositions, it stresses that heuristics are not absolute, but assessed with respect to a particular task and that even algorithms can be assessed in this way. By allowing students to generate their own solution paths, it helps make them conscious, creative members of the culture of problem-solving mathematicians. And, in enculturating through this activity, they acquire some of the culture's tools--a shared vocabulary and the means to discuss, reflect upon, evaluate, and validate community procedures in a collaborative process.”

Another approach that is complementary to the ‘epistemology of learning’ is the one of *constructivist theory*, whose roots can be traced back to the works of Piaget (Piaget, 1929,Piaget, 1970) and Bruner (Bruner, 1960,Bruner, 1966) . In general, this theory affirms that all knowledge is constructed, i.e. it is not the result of passive reception. This entails that learning is an active process in which learners construct new ideas or concepts based upon their current or past knowledge.

This approach stands in clear opposition to the one of a) *realist* (e.g., Platonic (Plato, 1981)) theories, for it does not support the idea that there is a "true" nature of things that can be successfully reached through a pre-defined method, and of b) *behaviorist* (e.g., Skinnerian (Skinner, 1965)) theories, since it tries to model the learning experience from the inside and not only describe it empirically from the outside. Within the constructivist framework, in fact, knowledge is not considered as a product but as a phenomenon that arises from the dialectic between the learner and what is learnt.

As often happens, over the years various currents have contested and modified the initial approach of Piaget by highlighting one or the other single feature as the most important. Without going into the details of such intellectual discussions, following a classification presented by Hein (Hein, 1991) we can enumerate some fundamental features of such an educational approach[1] :

I. Learning is an *active* process that requires the learner being engaged with the world.

II. There are always two different levels in the learning process: while constructing meaning, we also construct *systems of meaning*. For example, if we learn the chronology of dates of a series of historical events, we are simultaneously learning the meaning of a chronology.

III. *Language* has a central role in learning.

IV. Learning is a *social* activity: our learning is intimately associated with our connection with other human beings, our teachers, our peers, our family as well as casual acquaintances, including the people before us or next to us at the exhibit.

[1] The classification was originally created for a public of museum educators (notice the reference to an 'exhibit' in point IV), but it can apply to a large variety of educational contexts.

V. Learning is *contextual*: we do not learn isolated facts and theories in some abstract ethereal land of the mind separate from the rest of our lives: we learn in relationship to what else we know, what we believe, our prejudices and our fears.

VI. One needs *knowledge to learn*: it is not possible to assimilate new knowledge without having some structure developed from previous knowledge to build on. The more we know, the more we can learn.

VII. *Motivation* is a key component in learning. Unless we know "the reasons why", we may not be very involved in using the knowledge that may be instilled in us, even by the most severe and direct teaching.

In conclusion, for that regards our research work, this short review let us single out a number of key points that we believe should be ‘transformed’ into software requirements when designing an e-learning system. We are going to elaborate more on these concepts later, when giving an extensive definition of our approach.

Instead, in the next section we look more closely at a very specific learning context, namely, the interesting ‘situation’ of philosophy learners.

2.2.2 Learning Philosophy

The recent success of constructivist pedagogical theories have led many practitioners to agree on the basic fact that for a student to really understand something an active style of learning is necessary, in contrast with a passive reading and remembering of what is read. 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, for instance, is a very delicate matter: *philosophy*, as other subjects such as *theoretical physics*, *mathematics* and *logic*, deals only with abstractions. That is, in Laurillard terms, “descriptions of the world” (Laurillard, 1993). 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 (Kemerling, 1998,Mays, 1965,Kasachkoff, 2004).

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 (Carusi, 2003) that the three most important skills to develop in a philosophy student must be (a) analysis, (b) argument and (c) interpretation. It is worth reporting the original passage:

> Which skills receive most attention depends to some extent on the philosophical tradition in which teachers operate. UK departments of philosophy are predominantly analytical and so skills of analysis and argument will tend to come to the fore; a more historical and usually continental approach will instead tend to privilege interpretation and exegesis. However, in practice, the division is rather artificial, and I think that there can be substantial agreement that the three most important philosophical skills that we try to develop in our students are (1) analysis, (2) argument, (3) interpretation. In fact, these three skills are interwoven as analysis requires interpretation, and argument depends on the prior abilities to analyse and interpret correctly other philosophical positions.

In particular, in table 2-1 we detail Carusi's lengthier description of what each of the skills may entail, as far as the student is concerned.

SKILL	Description
Analysis	• analyse a philosophical problem or position into its component parts and be able to tell how they are connected together; • analyse an argument into premises and conclusions, and reconstruct the structure of the argument, filling in implicit premises where necessary; • analyse philosophical texts into sections and be able to see the connections between sections.

SKILL	Description
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.

Table 2-1. The three major philosophical skills (from Carusi, 2003)

In conclusion, although research demonstrates that we are quite far from a definitive explanation of the processes enabling the learning of philosophy (and similarly, of other 'academic' and 'abstract' subjects), we believe it is possible to build a software prototype which is also grounded on a pedagogical theory. In fact, following Carusi's analysis we consider the three skills above as a common denominator which lies at the heart of the various ongoing debates on the subject of teaching philosophy.

In the next chapters we will discuss in more detail how we intend to support the development of these skills by means of an ontology-based software environment (chapter 6).

Having discussed the learning theories informing our approach and also the specific epistemology framing the process of learning philosophical concepts, we must now mention a further characteristic of the learning process. That is, the fact that often we learn by *understanding and constructing stories*. The following paragraphs will clarify the importance of this specific type of learning, especially with regards to the implementation of our *pathways-inspired* software tool.

2.2.3 Learning through stories

Stories are part of our identity and our culture. Since we were born, we have been told stories that helped us making sense of ourselves, or of our role in the society; stories providing us with some understanding of the external world and a clearer perception of our life goals. According to Barthes (Barthes, 2000) narrative is present at all times, “in every age, in every place, in every society; it begins with the very history of mankind and there nowhere is nor has been a people without narrative” (p. 109). In the following paragraphs, in order to emphasize the connection existing between our understanding of the world and the ability to create and tell stories, we will briefly discuss the work of two influential authors, J. Bruner and R. Schank.

In general, the work of Bruner can be framed within the constructivist approaches mentioned in the previous section. Unsurprisingly, a major theme in his theoretical framework is that learning is an *active* process in which learners construct new ideas or concepts based upon their current/past knowledge. For example, this process may involve *selecting* information, *constructing* hypotheses and *making* decisions.

The key point here is that, in order to be successful, this process must rely on a cognitive *structure*. This cognitive structure (i.e. schema, mental models) is specifically what provides meaning and organization to our experiences - so that we can “go beyond” the information given and transform it into knowledge, i.e. integrate it into the already existing structures we possess.

In other words, learning is the continuous process of creating new mental models and adjusting them together with the old ones. Let us quote a passage from Bruner (Bruner, 1960) where he stresses the central the role of *structure* in the educational process:

> “The teaching and learning of structure, rather than simply the mastery of facts and techniques, is at the center of the classic problem of transfer... If earlier learning is to render later learning easier, it must do so by providing a general picture in terms of which the relations between things encountered earlier and later are made as clear as possible” (p. 12)

This is very related to another of Bruner’s ideas, the notion of a ‘spiral curriculum’. That is, a curriculum where the same concepts are re-iterated multiple times, so that eventually the student will be able to grasp the “full formal apparatus that goes with them” (ibid.: 13). According to Bruner, such a curriculum can be seen as a story-telling process, where the implicit *structure* of the stories (i.e. the interconnection of the events, the causal sequence of the actions) is precisely what guarantees the student’s learning.

This brings us to the other key idea we want to take from the work of Bruner, i.e. the conception of *stories*. Stories are the results of a *narrative way* of *knowing* (Bruner, 1996) :

> There appear to be two broad ways in which human beings organize and manage their knowledge of the world, indeed structure even their immediate experience: one seems more specialized for treating of physical "things," the other for treating of people and their plights. These are conventionally known as logical scientific thinking and narrative thinking . . . They have varied modes of expression in different cultures, which also cultivate them differently. [p. 39-40]

Narratives are therefore one of the preferred structures human beings use for making sense of reality. This thesis is further developed also in a more recent article (Bruner, 1991), where Bruner argues for a conception of the mind as constantly *structuring* reality using mediation through various "cultural products”. In particular, he focuses on the idea of *narrative* as one of these cultural products, providing a thorough description of ten of its characteristics.

In conclusion, the fundamental idea we want to reuse in our own work is the following: *the structure of narratives parallels the structure needed for successful learning*.

Let us now turn to another author whose work not only emphasizes Bruner’s views, but also elaborates on them: Roger Schank (Schank, 1990). His research spans from narrative theory to philosophy of knowledge, and constitutes a fundamental theoretical background when trying to grasp the relationships between human’s learning and storytelling abilities.

Essentially, he draws a line that connects intelligence, understanding, conversational structures and stories. According to Schank all of our knowledge scales down to the set of stories we are able to tell. Consequently the most interesting question becomes how we manage to get from one story to the other, namely how we constantly index new stories and relate them to the corpus of stories we stored in the past.

Within this approach, intelligence is defined as a “massive indexing and retrieval scheme” that brings out the linguistic representation of some latent conceptual structure. Schank offers a wide catalogue of the kind of stories we tell and are told, based on their *structure*, their *origin* and their *usage* and therefore states quite firmly the boundaries of what should be considered an intelligent behavior. For example, in the following passage he is analysing the set of stories ‘making up’ a military expert’s knowledge (Schank, 1996):

> "Knowing a great deal about a subject means being able to detect differences that will reflect themselves in differences in indexing. In other words, intelligence depends on clever indexing. Our expert is intelligent about military history. He sees nuances where others would not. He analyzes new stories well enough to be able to relate them to old stories that might not obviously be the same."

The *indexing of the stories* is thus what varies the most among different experts: in other words, stories can be seen as some sort of ‘raw’ material upon which we impose an organizing structure. As we will see in the following sections, structuralist theorists worked on a similar distinction, and called such structure the ‘discourse’. Also, in more recent times, specialized software application called *interactive narrative systems (Davenport and Murtaugh, 1997)* attempted to codify the structure (that is, the indexing of the stories) with the purpose of obtaining mechanisms for recollecting and navigating stories which resembled human's intelligent behaviors. We will say more about this in the following sections.

In conclusion, from our research point of view, Schank's work is particularly interesting because he is among the first researchers who attempted to use the learning process/narrative creation analogy as a model for artificial intelligence programs. In more practical terms, he set out to construct computer programs capable of *creating* and *understanding* stories, which rely on complex '*indexing schemes*' and '*composition rules*'. The same principle has been followed, more or less explicitly, by various other research enterprises. The narrative-inspired approach of our work (cfr. section 6.5.4) is, without any doubt, among such attempts. The scenario is quite a different one, though: the emergent availability of structured data sources on the Semantic Web, as we will see, poses new challenges to the traditional story-construction approach.

However, before getting to these topics, it is worthwhile spending a few words on the other strand of theoretical research that informs the 'storytelling' approach in artificial intelligence, i.e. the theory of narratology.

2.2.4 Narratology

Narratology is the name given to the critical and theoretical study of the numerous forms of narrative discourse, especially in literary and film studies.

If we want to briefly describe its theoretical origins, it is important to mention the influence of the work of the linguist De Saussure in contemporary narrative theory. In his book “Course in General Linguistic” (Saussure, 1995), published posthumously in 1916, the distinction between “signifier” and “signified” is discussed, and a general theory of language that claims the primacy of the *form* over the *content*, is presented. Basically, this means that the structure of the language, which can be abstracted from its everyday use, is actually the place where its essence lies. In other words, the semantic nature of a linguistic expression does not depend on its content, as apparently we are tempted to believe, but is instead tightened to the net of relations which constitutes the language phenomena as a whole. Another common way to express the same thesis is by saying that, in general, the *parts* obtain a meaning only *within* a wider *system*.

These ideas have been further developed by many scholars in different areas. To the aims of our investigation, it is worth noticing how they became central in the *French structuralism*. In fact, this theoretical movement through the 1960s and the 1970s has investigated the functioning of *narratives*, working on the assumption that, being linguistic phenomena, they could be studied by means of the methodology used by Saussure to analyze language.

Actually, an earlier appearance of narratology can be registered in the earlier intellectual movement named as Russian formalism. In particular, Vladimir Propp's “Morphology of the Folk Tale” (Propp, 1968) anticipated many of the methods of the structuralist narratological analysis in its breakdown of a corpus of Russian folk tales into a finite number of constituent parts. Precisely, thirtyone different morphological functions (mostly plot twists) and seven “spheres of action” (mostly characters) were identified.

In general, the basic idea of narratology is to scrutinize the internal relations of a narrative's component parts, and dissect how these relations are constructed in practically any given aspect of the narrative text (such as plot, narration, sequence of events, and so on). The text's *structure* can therefore be read as a system of meaning in its own right, which interacts with any apparent message the text contains. As a result, the concern of a narratological approach is not with what a narrative represents, but with *how it represents it*.

With reference to the work of the structuralists (McQuillan, 2000), but in particular of Genette (Genette, 1983) and later of Chatman (Chatman, 1978), we can sketch out the structure of a narrative as the union of:

- *Story*: it is the “what” of what is told, namely, the conceptual space representing people, events, objects and the organization of different entities. Somehow, it also refers to the abstract chronological structure of events. It corresponds to Saussure’s *signified*.
- *Discourse*: it is 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 this way, different effects can be created, such as humour or surprise. This category includes and is influenced by another one, the particular media used to deliver the narrative. In fact, the choice of the media will always affect the kind of rhetorical stiles allowed, or, for example, other time-related constraints. It corresponds to Saussure’s *signifier*.

- The *narration* itself: this dimension was firstly introduced by Genette, and refers to the unavoidable influence of the speaker on the final narrative's effect. It is the "Qui parle?" problem: every narration is always in a context, and therefore assumes some peculiar meanings from it, at least the point of view of the speaker.

This theoretical model, firstly developed to describe "stories" in the classic sense (namely novels, romances, or any other work in literature), has been extended to be used with any kind of media that can be possibly employed in the delivering of a narrative (myth, theatre, film or even hypermedia (Walker, 1999)).

2.2.5 Digital Narratives

The most interesting development of narrative theory, from this research point of view, is its translation into the digital world, namely, the existence of programs and languages to represent the dimensions linked to a narrative definition, support their dialectical interchange and, more generally, foster new ways to browse intelligently semantic spaces.

For example, in (Mulholland and Collins, 2002), tools based on narrative structures enable communities to celebrate and explore regional heritage. Communities are accustomed to discuss in electronic forums, where basically the main activity is the mutual exchange of stories representing a standpoint on a particular subject. These stories, if properly indexed (as we will see in section 2.4.1, through the use of *ontologies*) can be retrieved in novel manners, explored in a personalized way and compared using multiple viewpoints. In this way, the reader "does not just receive the narrative but actively constructs a story for themselves during the reading process".

This approach has implications on the learning side: learning is constructed and the navigation through the resources' description is at the same time autonomous (the user has control on it) and guided (by the semantic relations between the concepts-descriptors).

A more detailed explanation of the processes involved in the creation of a digital narrative is the one presented by Brooks (Brooks, 1996) (Brooks, 1997), one of the pioneers in this area. The aim of his research, as stated, is to employ the computer to "generate multiple narratives quickly and semi-autonomously" out of a pre-inserted "story" material.

This vision, in particular, applies to cinematic story construction, through the use of computer-based *storytelling systems*. According to Brooks, a storytelling system is not "a magic box which creatively makes up a story when asked, but a system of specially stored and organized narrative elements which the computer retrieves and assembles according to some expressed form of narration". In other words, such a system has 'knowledge' in order to create a *discourse* out of some specified content, and in doing so, it helps the author by providing an environment for non-linear, multiple point-of-view stories.

If normally the writing process produces a story that is then delivered to the audience, generating some feedback on the author, with the advent of the computer, claims Brooks, this process has changed. Computers, in fact, can provide some decisive support to the creation of stories, and not only as a word-processing tool.

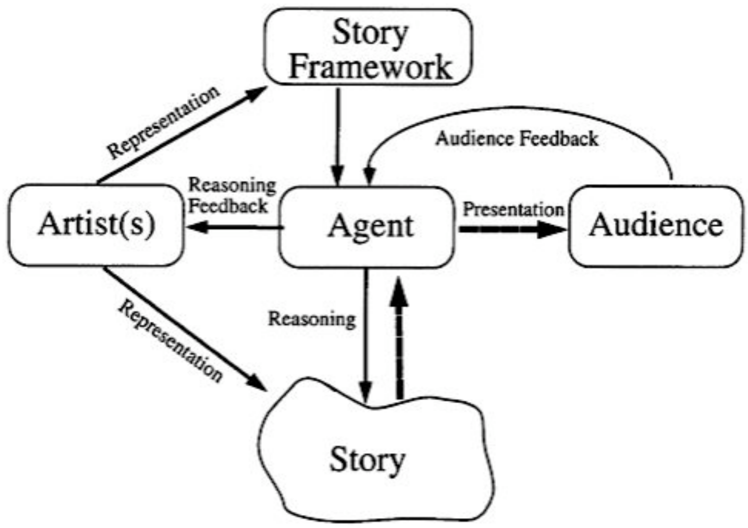

Figure 2-2 – The role of an agent in the story creation process (from Brooks,1996)

As we can see from figure 2-2, the whole process is focused around the notion of an *autonomous agent.* This is a software program that embodies the representation of a set of 'low level' competencies, each of them being the 'experts' in solving one small part of the larger problem domain. Thanks to these 'low level' representations the 'intelligence' of the agent supports a high degree of *adaptivity*.

In other words, this infrastructure lets an agent behave differently depending on the changing layout of the environment. Of course, the spectrum of possible behaviors an agent can exhibit is at the same time enabled and limited by the features (e.g., *number* or *granularity*) of the 'low level' components.

As already mentioned, the general framework and methodology to digitally represent a narrative has been instantiated in different areas. Among them, it is worth noting *game studies* (Ryan, 2001). In such a scenario, the interface between the user and the application is often centered on a *digital character* that acts in a virtual environment and triggers different story-paths, consistently with abstract story-structures encoded in the system.

A number of analogous applications exist in the *interactive digital-media* field of study. For example the Scene-Driver system (Wolff et al., 2004) is an educational software for children that, relying on the organization and description of contents adapted from an animated children's television series, aims at the reuse of the resources and at the involvement of the spectator on the development of a plot. In this case, the friendly user-interface is a domino-like board with tiles representing the different clips available. Children create sequences of tiles respecting some predefined rules (dependent on the narrative-compliant structure of the content), and subsequently can see a personalized animation in which they recognize a product of their sequencing of events and characters' actions.

In conclusion, following the analysis of Brooks (Brooks, 1996), we would like to underline two points which are of central importance to all the research in this area:

I. *Representation* and *reasoning* are inextricably and usefully intertwined: this means that at the core of a representation stands a conception of what constitutes intelligent reasoning; depending on the way we conceptualize a domain, we can present narratives which provide more or less added value and interesting features. For example, in the aforementioned Scene-Driver (Wolff et al., 2004) the authors intended to help children in constructing a meaningful sequence of cartoon-clips. Consequently, the underlying semantic representations needed to be describing various features of the cartoons' scenes (such as the name of actors appearing in a scene, or the list of the actions they are performing) as well as the more general 'sense' of the scene (e.g., 'conflict introduction', 'theme introduction' etc.).

II. *Granularity is a central problem in digital narratives*: granularity refers to the chosen unit size for building a story, and it embodies a trade-off between power and efficiency. If the story 'bricks' are smaller, there are more ways in which they can be composed together, but the representation and reasoning tasks may become incredibly complex and subtle; conversely, if they are larger, it is easier to put them together but there are less ways to construct meaningful stories out of them. For example, if a system aiming at the sharing of stories about a local community (such as Village Voice (Srinivasan, 1994)) were built using 'story bricks' modeled around the notion of *who is telling the story*, the resulting narratives would probably be trivial. Instead, by representing also the *events* told in the story, or the *physical objects* involved in the actions, the authors could support the construction of much more interesting (and less obvious) narratives.

As we will see, these two points are important also with regards to our work in the philosophical domain (see chapters 5 and 6).

We now have all the elements needed to start talking about the most recent developments in the World Wide Web, and how such developments can impact on the development of e-learning and narrative-oriented systems.

2.3 Semantic technologies for learning

The importance of e-Learning in the 21st century life has been repeatedly highlighted by Drucker (Drucker, 2000). He argues that the essence of e-Learning relies on the tools and knowledge needed to perform work being moved to the workers, wherever and whoever they are. This *just-in-time education* becomes therefore strictly integrated with the high velocity value chains that characterize nowadays commerce, and basically moves the focus of education from the institution to the individual.

Since various years, researchers have been looking at advanced techniques coming from the field of artificial intelligence in order to pursue further the idea of *just-in-time* and *personalized* education (Dillenbourg, 1994,Murray, 1999). In general, people refer to this research area as *artificial intelligence in education* (AIED).

For example, Dillenbourg (Dillenbourg, 1994) conceives computer based learning environments as systems representing open problem situations in which the learner can explore the consequences of his actions, and construct step by step his knowledge. An AI-enhanced e-Learning system, by taking advantage of *learner modeling techniques*, *pedagogical strategies* and *rule-based reasoning*, should assist the user in the decision making process and in the reasoning process. In this way, claims the author, it will eventually help turning the learner's declarative knowledge into operational skills.

Another example can be the work of Brusilovsky (Brusilovsky, 1996) on adaptive hypermedia systems (AH), technologies which draw their force from the construction of a learner model, a domain model and the definition of specific teaching behaviors depending on both of these models.

However, a downside of these systems is that (traditionally) they were created so to be used in a 'closed' environment. For example, the resources they work with are usually stored in a central repository, and univocally represented using a predefined semantic model.

Instead, as we have seen in the previous section, the emerging Semantic Web is characterized by a very different scenario: resources are usually distributed and, moreover, they can be 'described' using different semantic models. As a result, in more recent times researchers started investigating the specific implications of the scenario brought forward by the transposition of e-Learning systems to the Semantic Web (SWEL). Accordingly, these attempts will be also the main focus of our discussion.

In order to provide the reader with a better understanding of current research on SWEL, our review is organized as follows: first we look at five perspectives on SWEL systems, proposed in the literature (section 2.3.2).

Secondly, we examine different types of existing applications which employ these new technologies to support learning. In particular (following the order by which they are presented) these applications aim at:

a) Enhancing the learning-objects reusability by linking them to an ontological description of the domain, or, more generally, describe relevant dimension of the educational process in an ontology (section 2.3.3).
b) Providing a comprehensive authoring system to retrieve and organize Web material into a learning course (section 2.3.4).
c) Construct advanced strategies to present annotated resources to the user, in the form of browsing facilities, narrative generation and final rendering of a course (sections 2.4.1, 2.4.2, 2.4.3 and 2.4.4).

However, before starting a detailed review of SWEL system, it is worthwhile spending a few words on the Semantic Web vision and on the important characteristics of its technologies. By doing so, it will then be easier to understand how the SW is influencing the e-learning area.

2.3.1 The Semantic Web

The Semantic Web (SW) is described by Berners-Lee (Berners-Lee, 1998,Berners-Lee, 1999,Berners-Lee et al., 2001) as an “extension of the current web, in which information is given well-defined meaning, better enabling computers and people to work in cooperation". Basically it is an attempt to add a layer of *machine-processable* data to the existing web, so that software agents could carry out sophisticated tasks which would otherwise not be possible.

This 'semantic layer' can be composed of *metadata* (literally, data describing other data) specifying, for example, the intended meaning of already existing web-pages, but not only. In fact, much of the semantic layer could simply describe the *world* itself, e.g., think about some geo-political data about a country, or data representing a person's culinary preferences.

The only condition that all data on the 'semantic layer' must satisfy is that they should be represented using a *formal language* which can be 'understood' by the other computers in the network. By doing so, agents (software and human) will be able to unambiguously *query* the Semantic Web, as if it were a giant database, thus allowing systems to carry out tasks relying on highly distributed information sources.

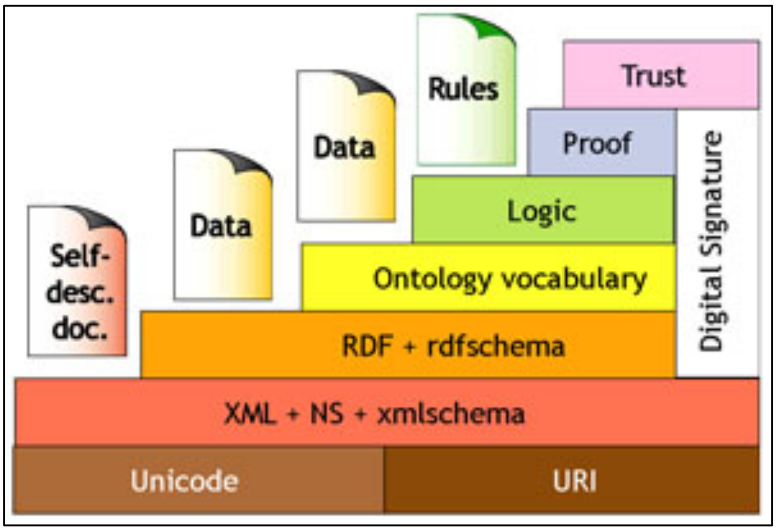

Figure 2-3 – The semantic web layers (from Berners-Lee 1999)

First of all, in order for this to happen, data must be *structured* and related to sets of *inference rules*. Computers could therefore conduct some automated reasoning on this huge knowledge base, as long as it is formalized following some consistent and well-known techniques.

This framework, as shown in figure 2-3, is usually characterized as consisting of a series of layers. Each layer represents a different technology, which takes advantage of the representational power of the technology underneath, to provide an abstraction capability to the technology above. So, for example, while the XML layer represents the structure of data, the RDF layer represents the meaning of data; the Ontology layer, instead, represents the formal common agreement about meaning of data; above all these stands the Logic layer, which enables intelligent reasoning over meaningful data.

Among these technologies, especially two of them are already widely used and are constantly augmenting the semantic mark-up of the old Web pages: the eXtensible Markup Language (XML) and the Resource Description Framework (RDF) (Decker et al., 2000).

XML lets users define their own tags, and use them within normal Web pages in order to provide arbitrary structure to a document: in this way the content of a resource should appear clearly and, if the semantics of the tags is known, a software parser could analyze the page and get specific results (that eventually will always point at some URIs). RDF, instead, aims at expressing the meaning of the resources, describing them through a *subject-verb-object* codification style. The triples thus obtained could be written using XML tags. Their semantics essentially says that a particular thing 'X' has a property 'Y' with a value 'Z'. In this way, the triples create *webs of information about related things*, and connect resources in a more human-like manner.

Even if the RDF layer already offers quite a strong representational power, it is not enough in order to overcome many ambiguity problems: for example, two different words (or identifiers) could be used in order to refer to the same concept. Consequently, there is the need of a meta-level that describes these common meanings, something like a document or a file structure defining mappings between different databases. Ontologies (Corcho and Gomez-Perez, 2000,Gruber, 1993,Noy and McGuinness, 2001) provide this crucial functionality, and are therefore a key area of research.

There are various views around the definition of an *ontology*. Here we use the definition of Gruber of an ontology as an "explicit specification of a conceptualization". In other words it is a formalized theory of what exists in a particular domain. The simplest ontology is a *taxonomy* (namely, a tree-like data structure that defines classes of objects and relations between them), endowed with a set of *inference rules*, which allows advanced manipulation of the classes (for example, some cross reasoning between them). It is a widely accepted tenet in SW communities that ontologies constitute the *backbone* of the Semantic Web, since their expressiveness transforms them into some sort of *universe of discourse* for data manipulation.

Within such a scenario, a number of different applications and systems that make use of the technologies presented above have been developed. It is not the purpose of this review to go through all these attempts, so we will just recall some examples that deal with *browsing*, *retrieval* and *classification* of resources.

Let us start by recalling the classic keyword search performed through an engine like Google (Brin and Page, 1998). This technology finds us a list of resources that, although being retrieved and ordered through a variety of mechanisms (ranking algorithm, string similarity algorithm etc.), is essentially the result of operations performed on the *syntactic features* of the web documents (hyperlinks, strings). The primary consequence of this is that we may get unwanted results. Being a little more demanding, we cannot have a 'contextualized view' on the result set i.e. a map showing what the results *mean* in relation to each other. Such a map could help us in carrying out common higher level tasks (e.g., *reviewing*, *analyzing*, *studying*).

The situation just described would not be happening, instead, when using 'semantic' search. That is, if we considered the metadata associated with the resources we could get a series of results from different repositories, already organized according to some ontological knowledge. In this case, results are gathered and presented on the basis of the *semantic description* of the web documents (i.e. on one of their declared meanings).

A seminal system like this is the one described by Guha and colleagues' (Guha et al., 2003). Building on the important "distributed extensibility" characteristic of the SW, they propose an improvement to the traditional Information Retrieval (IR) technologies, which, claim the authors, are based almost purely on the occurrence of words in documents. Essentially, their system (called TAP) is an infrastructure that provides simple mechanisms for sites to publish data onto the SW via a minimalist query interface called *GetData*. Pure HTML pages are *scraped* and knowledge is formalized into RDF files. This technique, plus the manual annotation of other resources, has supported the creation of a knowledge base which is used for testing various ‘semantic’ search mechanisms. These results are then formatted and presented to the user as augmented ‘classic search’ results.

Another classic example of a SW enhanced system is the Simple HTML Ontology Extensions (SHOE) language (Heflin and Hendler, 2000), an application of SGML and XML that allows users to define extensible vocabularies (that means, ontologies) and associate machine understandable data to them. In this way, Web pages can be easily marked up and searched using the ontologies. For example, the authors describe an interface where ontologies are selected from a drop-down menu, effectively providing a context for the search through the usage of their knowledge encoded in classes and relations.

Finally, Piggy Bank (Huynh et al., 2005) is a more recent example worth mentioning, for it integrates the benefits of SW technologies within the more familiar web-browser. More precisely, PiggyBank tries to augment the users’ experience of the Web by giving them the possibility to extract individual information items from within Web pages and save them in SW format, that is, encode them as metadata. This approach, claims the authors, aims at resolving one fundamental problem of the SW, the scarcity of annotated resources. This is described by the authors as a ‘chicken-egg’ problem, which could be solved with the integration of a SW tool into a web browser.

This ‘minimalist’ review of SW-oriented applications gives the reader a first grasp of the technology’s intended usage. At the time of writing, there is already a quite vast (and constantly growing) number of research and industrial projects investigating the adoption of SW in various domains. For an extensive review of the potential applications and research directions we refer the reader to other publications (Antoniou et al., 2008,Passin, 2004).

It is now possible to focus our attention on the applications of the SW in a very specific area, computer supported education.

2.3.2 Semantic Web for e-Learning: roadmaps

Since the employment of SW technologies in e-learning is a quite new research area, many are the authors who directed their efforts towards the precise definition of the *broader scenario* within which such technologies could be used. In fact, quite often such scenario-definitions are used by other people as *roadmaps* pointing to the various possible research directions. In this section we review five of these roadmaps.

A seminal work in the emerging Educational Semantic Web was done by Mizoguchi and Bordeau (Mizoguchi and Bourdeau, 2000), who defined the new

'Instructional Design' paradigm as the evolution of Intelligent Tutoring Systems and Interactive Learning Environments, that is, as a "process by which learning events can be defined or described, independently of their instructivist or constructivist orientation". This new paradigm is fostered by the introduction of ontological engineering in the educational field: in fact, thanks to the precise semantic representations of an ontology it is possible to map out the wide range of existing solutions to the most common problems of instruction.

An ontology, claim the authors, can be described as a three levels device: at the first one it appears as a structured collection of terms, a *taxonomy* that elicits the concepts' hierarchy in a particular domain; at the second level it provides formal definitions of the concepts, relations, constraints and axioms, all of which make the ontology more operable for computer agents; at the third level, the ontology is *executable* in the sense that "models built based on the ontology run using modules provided by some of the abstract codes associated with concepts in the ontology".

If this view is translated into the educational scenario, according to the authors, the first feature supports the *sharing* of domain conceptualizations between humans, as a common vocabulary for representing the knowledge; the second feature enhances *computer's intelligence*, and therefore bridges gaps between humans and computers; the third one, instead, makes this knowledge *operative* and let computers decide actions to perform within a system thanks to activity-related concepts (especially in *task* ontologies (Chen et al., 1998)).

Stojanovic and others (Stojanovic et al., 2001) agree about the fact that ontologies are the most important improvement the SW brings to e-learning technologies. In an ‘Educational Semantic Web’ the everyday activities would be ontology development, ontology-based annotation of learning materials, composition of resources in learning courses and active delivery of the learning materials through e-learning portals. The authors define three principal criteria (see fig. 2-4) for locating learning materials, namely, what the learning material is about (content), in which form the topic is presented (context) and how it is presented in relation to other materials in a learning course (structure).

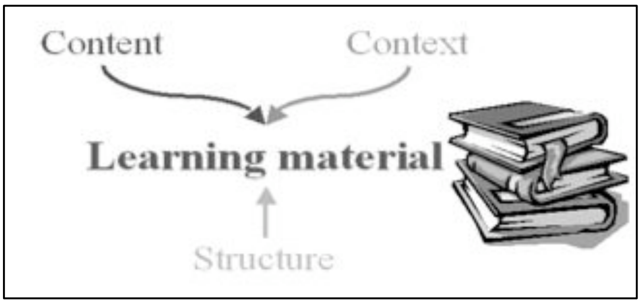

Figure 2-4 - The learning dimensions (from Stojanovic, 2001)

Following this classification, they provide some examples of the kind of ontologies that could support the description of a learning resource. *Domain* ontologies, for that regards the content, would solve problems due to language ambiguities, and would evolve basic keyword queries into semantic searches. A *context* ontology, instead, would identify learning contexts such as an *introduction*, an *analysis* of a topic, or a *discussion*, or presentation contexts such as an *example* or a *figure*. Finally, *structure* ontologies would serve to specify the construction-grammar to assemble small bits of information into personalized and quick-delivered learning narratives; concepts like *Prev, Next, References, IsBasedOn* etc. constitute the semantic connections to build a “Lego” learning system tailored to meet individual skill gaps.

The three dimensions, obviously, would also provide the main pathways to access a learning repository: resources can be accessed through a semantic query on one or more of them, or through a conceptual navigation based on the ontological representations available.

Another overview of the future implications of ontology usage in teaching and learning is proposed by Wilson (Wilson, 2004), who summarizes the potential benefits as follows:

- Students are provided with advanced browsing and searching support in their quest for relevant material on the Web.
- Syntactically different but semantically similar resources can more easily be located.
- The same work involved in creating an ontology can directly benefit learners by helping them to visualize and comprehend the relationships between concepts in their domain.
- Information can be shared across educational applications, enabling reuse not only of learning objects but also of domain knowledge and pedagogical strategies.
- Learners can be provided with the intelligent and personalized support that they would otherwise miss out (for example, personalized courses can be generated on demand).

In a similar way, the author outlines also the implicit risks of a serious employment of the technology in the educational areas:

- The ontology development process can be difficult and costly: the more expressive the ontology, the more complex and time-consuming this task; moreover, achieving an 'objective' representation of a domain is next to impossible.
- The context within which an ontology is supposed to be used tacitly constraints the definition of its concepts; so, for knowledge to be effectively shared, this contextual information must be formalized as well.
- Rich and complicated ontologies carry great expressive power, but are hard to comprehend especially for end-users.
- Since communities from different backgrounds (like library science, knowledge engineering, business) are involved in the ontology development process, there is a lot of overlap and reinvention, or many cases where the same things are defined differently.

According to Devedzic (Devedzic, 2004), the adoption of SW technologies can lead to an improvement in AIED (Artificial Intelligence in Education) only if the new technologies are firstly properly understood and digested.

In fact, various characteristics of the traditional ITSs (Intelligent Tutoring Systems), for example ELM-ART (Weber and Brusilovsky, 2001)), are already grasping key aspects of the learning experience. Thus the problem is to determine exactly where the ontological framework fits best, and how to use it (annotation of resources, or just representation of usable knowledge).

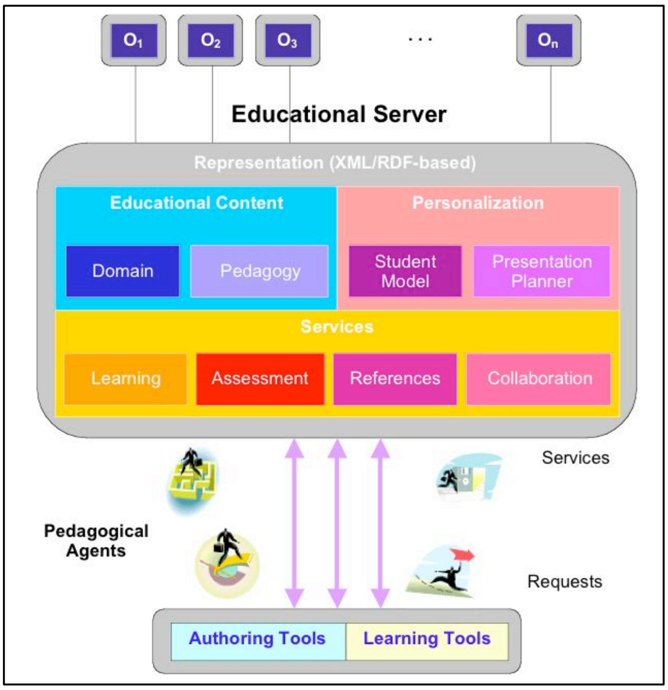

Figure 2-5 – Schema of a Semantic Web Educational Server (from Devedzic, 2004)

The model he presents (see fig. 2-5) is very useful for it takes into consideration different SW technologies and all the possible protagonists and scenarios involved in any learning activity. We can briefly summarize its main features with the two following points:

- Ontologies are the backbone of the system: they are used to codify different levels of shared understanding, like the vocabulary, the semantic interconnections, rules of inference, and to provide the structure used to semantically markup the resources available (this markup is then recorded in other formats, like XML, for better interoperability). The kind of ontologies needed to cover the whole learning experience should be about *domain* characteristics, *pedagogical* approaches, *student* models, and *presentation* styles.

- Services like search agents, information brokers, filters and integrators constitute the interface between the users and the knowledge base of the system. Moreover, they guarantee also interoperability between different applications on the Web at the semantic level, allowing the end user to be employed in complicated operations of *learning* (course offering, integration of educational material, tutoring, presentation), *assessment* (on-line tests, performance tracking, grading), *reference* (browsing, search, portals) and *collaboration* (group formation and matching, class monitoring).

Finally, let us conclude this section by mentioning the work of Stutt and colleagues (Stutt et al., 2005,Stutt and Motta, 2004). The researchers describe in a detailed way a scenario where one of the major problems of the SW, the competing and overlapping nature of its ontologies, would be overcome by the existence of a multiplicity of community-based Semantic Learning Webs (SLWs).

In fact, since the nature of the medium is distributed, it makes sense to let agents construct ontologies and repositories in a distributed way. Communities would build so-called “knowledge charts” (see fig. 2-6), in order to represent the information of their interest, while specific “knowledge browsers” would navigate this digital spaces looking for consistency and correlation between concepts.

The issue the authors address is essentially the need for *context* in these learning process. In fact, relying on various communities and not on a central and ‘objective’ repository, the technology offered by the SW could support one fundamental learner’s necessity: the possibility of structuring and locating a single piece of knowledge within a local panorama (the knowledge chart), and possibly, be able to move on to related areas (other neighboring knowledge charts). The *interpretation* of information is thereby fostered by the navigational capabilities of the SLWs.

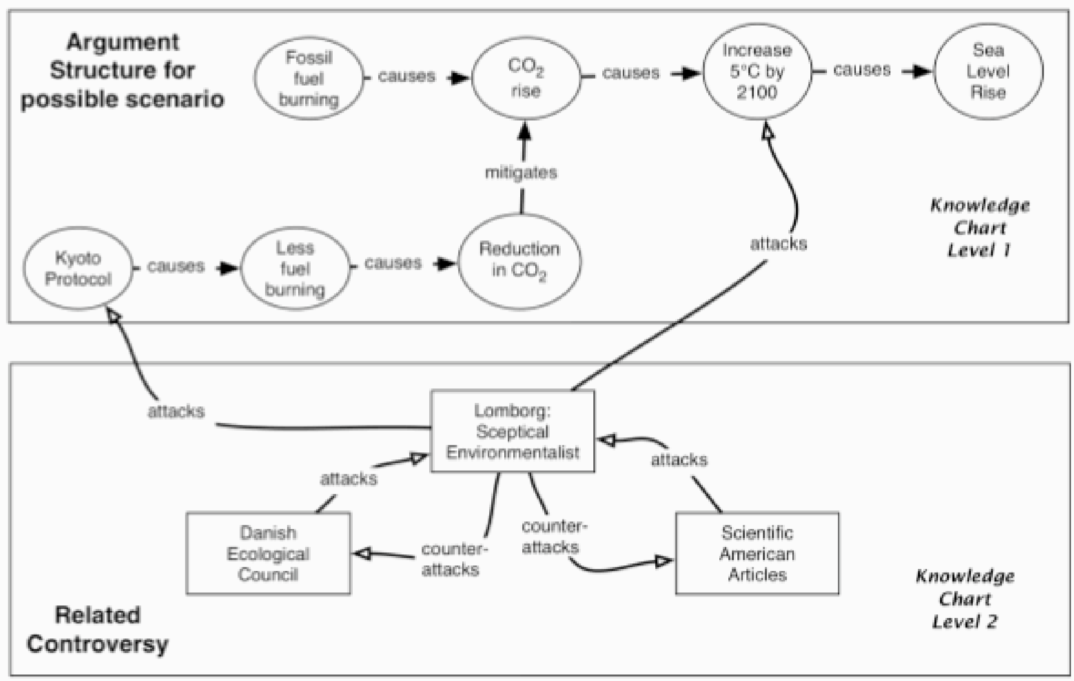

Figure 2-6 - Example of knowledge charts related to Global Warming (from Stutt, 2005)

After this summary of SWEL's generic frameworks, we can now examine some examples of real systems that make use of ontological engineering (and other SW technologies) to overcome various problems connected to the traditional learning technologies.

2.3.3 SWEL examples (I): ontology-enhanced e-learning

With the label ‘ontology-enhanced e-learning’ we want to refer to a class of research tools attempting to solve the common problems associated with *learning objects’* usage. In order to understand better the nature of these problems, we should proceed one step back so to briefly explain what are learning objects (LO).

The distributed nature of the Web, and consequently the lack of central organized repositories for digital resources, has led to the creation of *descriptors* in order to foster exchange and re-use between these resources. In the case of educational resources, the notion of ‘learning object’ (Duval and Hodgins, 2003) has been developed in order to frame the basic independent units usable in a learning activity. A learning object is thus defined as any entity (digital or non-digital) that may be used for learning, education or training. LOs are normally composed by a content and a set of descriptors. These last ones, usually called metadata, should apply to LOs in order to describe their salient features, and facilitate their exchange.

During the years, many standards have been proposed for defining LOs schema. Among them, let us mention the Dublin Core Metadata Initiative (DCMI) (DCMI, 2008), the IEEE Learning Object Metadata (LOM, see fig. 2-7) (IEEE), the ADL Sharable Content Object Reference Model (SCORM) (ADL, 2004) etc.

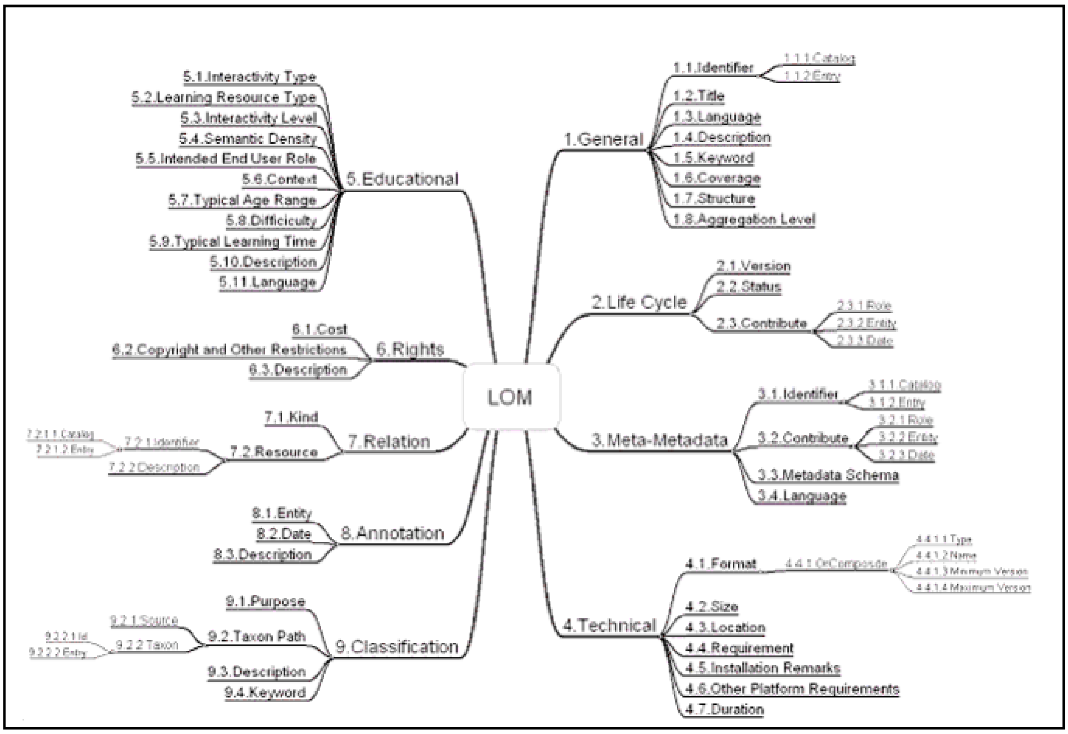

Figure 2-7 - The LOM metadata schema (from http://www.imsglobal.org)

These are just the most important attempts to create metadata standards in order to facilitate communication and re-use of learning resources. In fact, quite a few others are available, with the consequence that sometimes the various metadata schemas are not compatible with one another. A key limitation of LOs is that *their metadata are still semantically poor*. The information they provide remains superficial, proper for a simple course construction or for the location of resources through key words, but not capable of linking them in a non-sequential manner. So, the features that carry the biggest pedagogical advantages and implications, such as advanced semantic browsing facilities between educational resources, are extremely limited.

Therefore, various researchers have supported the usage of Semantic Web languages (in particular, ontologies) as the main strategy for surmounting the restrictions imposed by LOs metadata schemas. For example, the detailed definition of the salient features of a knowledge domain through a semantic language can be used to map different learning resources; also, this could allow a *meta level* of reasoning able to produce personalized and pedagogical narrative structures.

The augmentation of a LO can be achieved along different dimensions. The ontologies related to this task, in fact, try to model not only the *domain* knowledge of the learning material, but, more generally, all the aspects possibly involved in the educational experience. We will now discuss some existing research showing what are these aspects and what kinds of improvements are thus achieved.

The work of Koper and his team (Koper, 2001,Koper and Olivier, 2004) at the Open University of the Netherlands is a fundamental milestone in the field of advanced e-Learning technologies. Quite importantly, one of the starting points of their research is the recognition that the success of an e-Learning strategy does not rely on the medium (Internet) itself, but on the pedagogical design used in conjunction with the medium. Therefore it comes natural to formalize and clarify the dimensions of a pedagogical design, in order to instantiate them during every learning event.

A *unit of learning* is defined as the smallest meaningful ‘chunk’ of a learning event, that is, the smallest building block capable of carrying its own meaning and effectiveness towards the attainment of a learning objective. Learning objects, being extremely poor in their metadata definition, cannot fully exploit their most important feature, namely their being re-usable educational entities. In order to make their usage as flexible as learning management systems would like it to be, the authors have leveraged the difference between learning objects and units of learning by defining precisely the pedagogical dimensions of the latter, and use this meta-model to exchange and work with the formers. In other words, as they say, their basic idea is to:

I. Classify, or type, the learning objects in a semantic network, derived from a pedagogical meta-model,

II. Build a containing framework expressing the relationships between the typed learning objects and

III. Define the structure for the content and behavior of the different types of learning objects.

An Educational Modelling Language (EML) (EML, 2005) has been defined in order to describe the features of a unit of learning. These features basically represent the *meta-model* behind any pedagogical model, that is, an abstraction at the same time capable of expressing semantic relationships between pedagogical entities and of remaining pedagogical neutral. The meta-model is composed by four packages (figure 2-8):

1. The *learning model*, It describes how learners learn based on accepted consensus among learning theories. There are concepts like *external world*, *situation*, *cognitive state*, *stimulation*.
2. The *unit of learning model*. It describes how real instantiations of learning practices are created, given the learning model and the

instruction model. Basically it contains the knowledge necessary for designing a learning event. It deals with issues like the *roles* of *staff* and *learners*, the *objectives* of a *group*, the *prerequisites* of the *learners*, *context* and *assessment* of learning etc.

3. The *domain model*. It gives information about the type of content and the organization of that content. In fact every content domain has its own structuring of knowledge, skills and competencies (e.g., math, or philosophy)
4. *Theories of learning and instruction*. It formalises the theories present in the literature, and collects them into four categories. The *empiricists*, adopt a purely behaviouristic approach. They assume that knowledge is based on experience and that processes can be observed, predicted and analyses independently of the context and of the internal state of the learner. The *rationalists*, focus on cognition as the medium between a person and the environment, and therefore treat it as the real force that generates knowledge. The student is given a central role in the education process, since he is the builder of his own knowledge. The *pragmatic and cultural historic* approach, instead, considers knowledge as distributed between individuals, tools and communities, thus locates in the situation and the cultural-historical context the determining forces that drive the learning experience. At last, the *eclectic* model, combines different features from the other three positions.

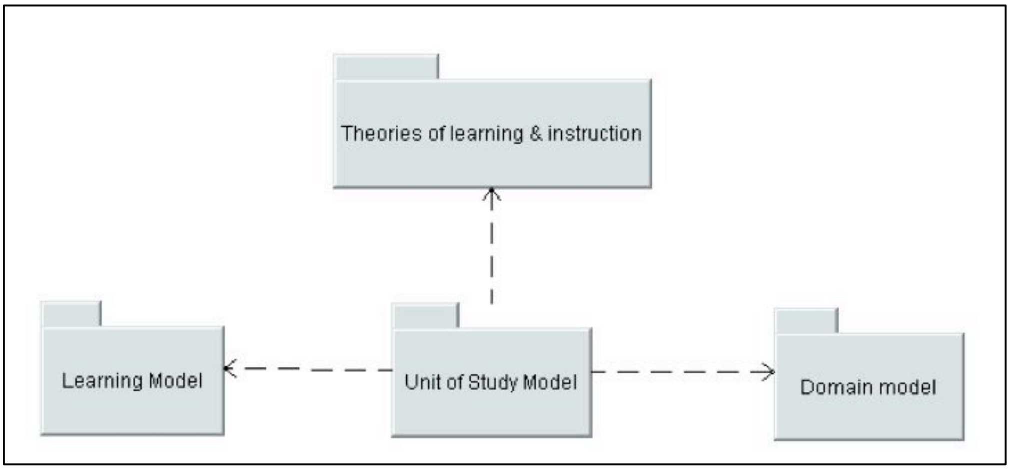

Figure 2-8 – The dimensions of the pedagogic meta-model (from Koper, 2001)

The integrated meta-model should therefore overcome the LOs shortcomings by explicitly declaring the fundamental constraints of any educational activity. The model is further analysed by the authors along seven basic requirements (completeness, pedagogical expressiveness, personalization, compatibility, reusability, formalization, reproducibility) and judged capable of enhancing what can be done in online learning. Some of the expected outcomes are:

- Coordination of multiple users.
- Integration of learning objects and services.
- Providing a learning activity layer over learning objects and services.
- Supporting dynamic personalization/adaptation.
- Supporting multiple pedagogical approaches.

Gasevic and colleagues (Gasevic et al., 2004) propose a system that improves LOs usability through domain ontologies. They start by highlighting that, due to the dual structure of LOs (metadata plus content), there could be two usages of ontologies in regards to them: ontologies that describe LOs' metadata, and ontologies that describe LOs' content.

The first ones act on top of standard metadata schemas, like the ones we have introduced above (e.g., LOM), enriching their meaning and giving more context to their usage. The second ones instead are domain ontologies which authors can create and use in order to semantically mark-up directly the content of a learning resource. Later, the teacher can extract annotated parts of documents and re-assemble them into a presentation or a course.

This second solution is the one the authors point out as the key advantage of SW technologies on e-learning, supporting the maximum reusability and semantic "freedom" (since different referring ontologies allow different semantic mark-up on the same document). The figure below illustrates the system's workflow and the technologies involved.

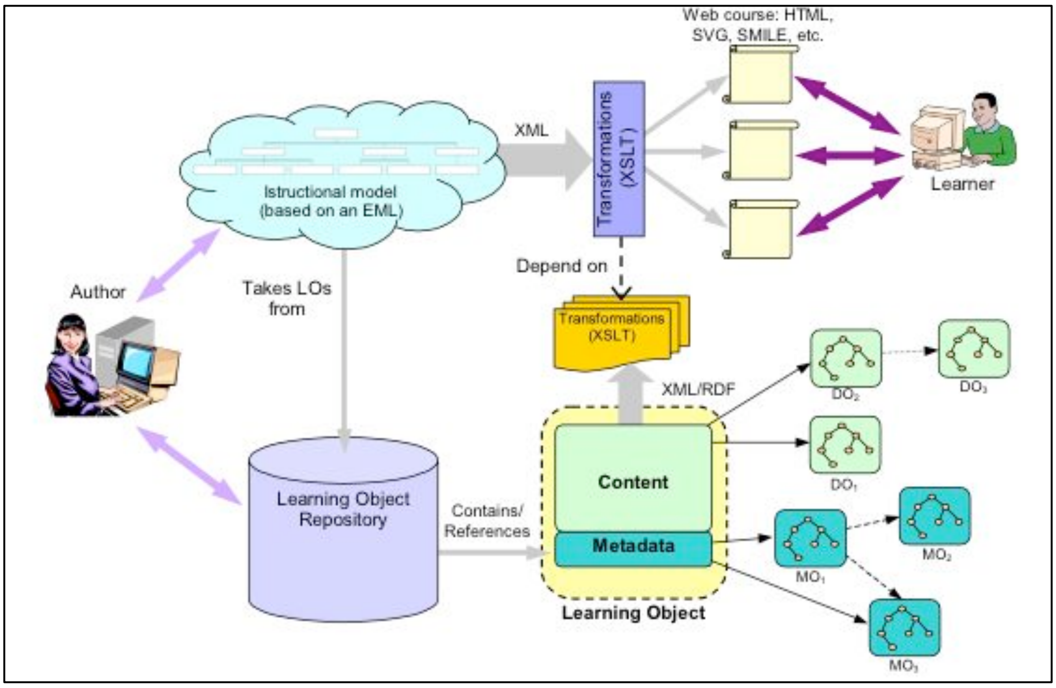

Figure 2-9 – eLearning process enhanced by ontologies (from Gasevic, 2004)

It is worth noting a few characteristics of such architecture (see fig. 2-9):

- LO's content can be produced in many different ways (text, slides, video, etc.), but its description should be encoded into some well-known metadata schema. LOs repositories are distributed sources of LOs and can contain either the metadata associated with them or the reference to the metadata on the Web. An author accesses LOs and integrates them into an instructional model of a course, designed according to an education modelling language (e.g., EML).
- Ontologies are used for both the description of metadata (MO) and the description of the content (DO). These ontologies do not necessarily have to be created by the authors, although some user-friendly tools can help them build their knowledge models, for example, during the annotation phase.
- The annotation process relies on the ontologies available, and is supported by specific tools capable of producing semantically marked up Web resources (that in this case are LOs) from different raw documents, like for example HTML pages, Scalable Vector Graphics (SVG), etc.
- The communication between the different technologies is guaranteed by the usage of XML and XSLT. The first one can be easily obtained from an EML instructional model, or from LO's annotations, through common exporting features. The second one, eXtensible Stylesheet Language Transformation, automatically transforms the XML structured data into a learner-suitable presentation, formatted in basic HTML.

Elena (Simon et al., 2003) is an analogous application that addresses the limitations of metadata standards through the usage of ontologies and P2P technologies. Elena is defined as a *smart learning space*, that is, a mediation infrastructure for Educational Services (ESs). Since LOs do not provide enough vocabulary to model a real course they should be replaced by educational services, whose data model, instead, keeps into account the *pedagogical context* in which the service is offered (it describes things like *educators*, *resource type*, *technology type*, *physical places*, *terms* and *conditions*, *schedule*).

A personal learning assistant (PLA) should therefore be able to query ESs in order to collect resources matching personal profiles and specific learning contexts, using some common technologies associated with Web-services (SOAP, WSDL, DAML-S).

The broader framework of this scenario is a smart learning space, namely an infrastructure based on Edutella (a peer-to-peer technology to connect highly heterogeneous educational repositories (Nejdl et al., 2002)) within which ESs can operate and collect resources or information about resources. This learning management network is queried by the PLA, taking also advantage of the learner's profile in order to personalize query results.

This system adopts ontological technology in order to overcome different problems:

a) To enable the various actors (educational services providers, that is, the peers in a P2P network) in the smart learning space to communicate with each other on a high level of abstraction.

b) To represent the natural language query within a semantic formalization: for example, the query "find a tutorial that explains the semantic Web to a novice" would make use, at least, of an ontology of learning resources ("tutorial"), an ontology of computer science ("semantic Web"), an ontology of learner's profiles ("novice") and an additional ontology that describes Web services' capabilities and query methods.

c) To annotate the learning resources.

The retrieval of the desired educational materials many times is not only a matching procedure between constraints and annotations, but exploits also other SW technologies (like RDF), locating resources with the help of reasoning engines.

The three examples just presented are frameworks aiming at improving the capabilities of the traditional Learning Management Systems by means of ontological engineering techniques. Before moving on to the next group of SWEL systems, it is worth noting that much research in this area has also been specifically devoted to the creation of 'educational' ontologies, i.e. ontologies which can then be used as *components* in more sophisticated frameworks.

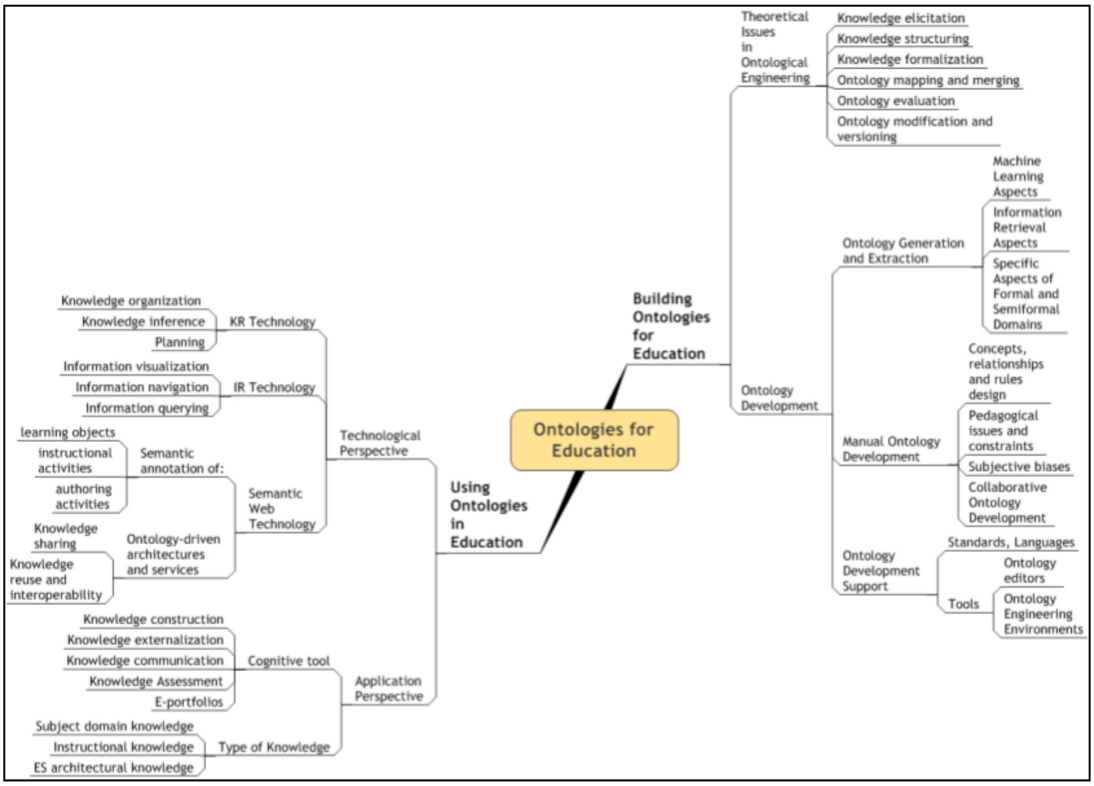

Fig 2-10 - Ontology of ontological technologies for education (from Dicheva, 2005)

So, for example, it is possible to find:

- ontologies trying to capture the ‘essence’ of a learning resource, namely, its function abstracted from the particular domain it refers to (Ullrich, 2004);
- ontologies abstracting the concepts useful when deploying a learning design, including notions describing learners’ processes (Lama et al., 2005), (Mizoguchi and Bourdeau, 2007);
- ontologies of the goals of IT in education (Kasai et al., 2004);
- ontologies for authoring intelligent educational systems (Aroyo and Mizoguchi, 2004);
- ontologies modeling learners’ profiles and behaviors (Chen and Mizoguchi, 2004);

- ontologies aimed at supporting collaborative learning environments (Inaba et al., 2001) (Barros et al., 2002)
- 'self-referential' ontologies mapping out the wide spectrum of SWEL enterprises (see fig. 2-10), so to help in creating portals gathering the diverse research attempts (Dicheva et al., 2005) .

2.3.4 SWEL examples (II): authoring systems

A more complex type of application in the SWEL are the systems that aim at supporting an author of a learning course in collecting resources and organizing them in relation to a particular approach or point of view. In this case, ontologies are used not only to describe a pedagogical strategy or the structure of a domain, but they become the main instrument to avoid "re-inventing the wheel" every time, and to tackle the exponential growth of courseware and learning materials. The systems presented below, therefore, simultaneously achieve different goals: e.g., they produce reusable courseware, emphasize the structure and the modularization of the authoring process, and keep into account existing standards to support interoperability.

The Courseware Watchdog (Tane et al., 2004,Tane et al., 2003) is an ontology-based application built at the University of Karlsruhe, in order to tackle problems such as the increasing number of topics in education and the decentralization of resources on the Web. The retrieval, interaction and management of resources is becoming increasingly difficult, so, say the authors, a new and more comprehensive approach which integrates the content, structure and evolution of the courseware material is needed. In their words, “true interoperability does not only need data integration, it also has to consider the integration of applications”.

In order to reach this aim, the Courseware Watchdog, building upon the KAON framework (KAON, 2004) (an open source ontology management infrastructure, compounded of tools for ontology management and application), is deployed as a set of tools (see fig. 2-11) that assist the user in:

a) Understanding and browsing ontologies,
b) Retrieving relevant material
c) Querying semantically annotated resources repositories,
d) Organizing the collected documents,
e) Updating the ontology.

The broad range of activities supported by the application strongly relies on the usage of a single semantic model (the KAON framework), and it is worth describing it more precisely: .

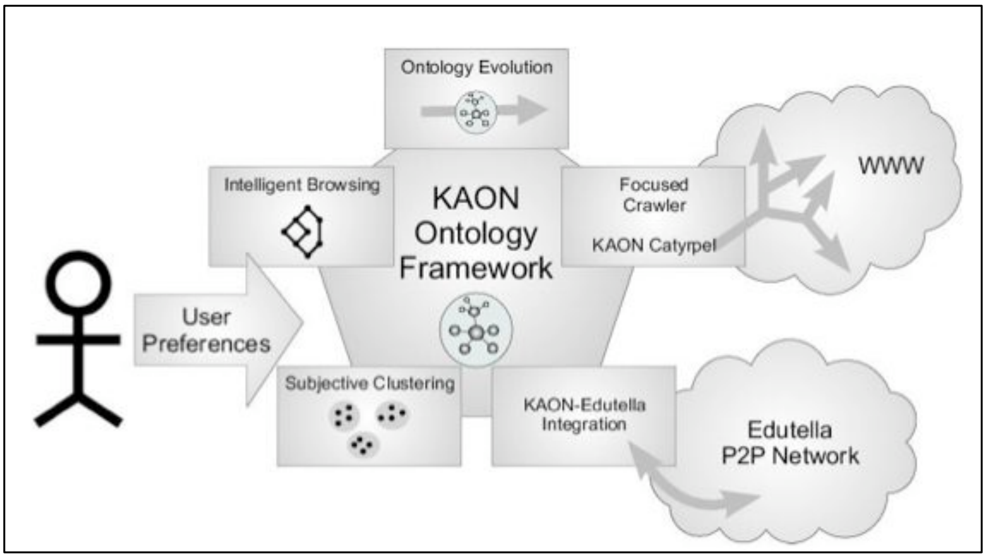

Figure 2-11 - The components of a Courseware Watchdog (from Tane, 2003)

a) The user may or may not use an ontology he has created; therefore, it is vital for him to familiarize with the concepts and relations employed. Visualization techniques (display of hierarchies through concept lattices, in particular, using Formal Concept Analysis) and browsing techniques (relational browsing) are used in order to improve the interaction between the user and the content. Moreover, the visual interface is conceived in order to maintain always an open perspective on the ontology, as it were a map to browse (see fig. 2-12): the left side of the screen shows it all the time, while the right side changes depending on which of the other tools is currently active.

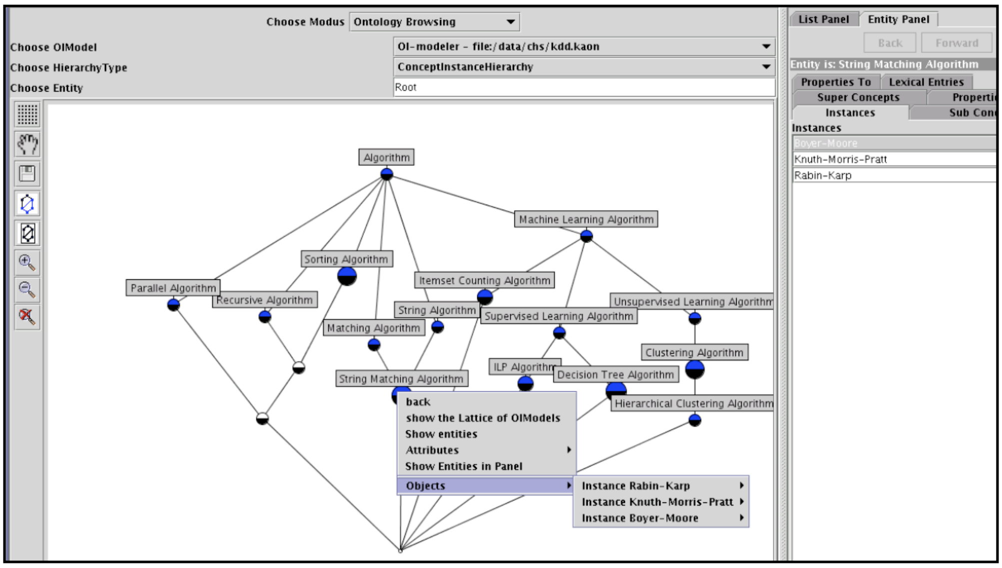

Figure 2-12 - The browser component of the Courseware Watchdog (from Tane, 2003)

b) A “focused” crawler (i.e. a program that collects data from the web automatically by following links extracted from web documents) uses the ontology to direct its research, and lets the user define some preferences, like the weight to assign to different concepts or relations, or how large a radius around a selected entity is to be considered. This allows different levels of "sharpness" in the focusing, and the recording of the retrieved results (pages crawled, position of the relevant entities, link structures between pages) in a knowledge base.

c) The Courseware Watchdog can also access in the Edutella network (Nilsson, 2003) as a full working peer, capable of querying for metadata on learning objects and of publishing local resources in the network. Since Edutella works mainly with RDF metadata, specific APIs have been developed to guarantee the communication with KAON ontology language and data model. The user is given an extensible set of basic query templates that can be filled in and employed straightaway, in order to facilitate inexperienced users to pose meaningful queries.

d) The data retrieved are stored as instances in the knowledge base of the relevant ontology, therefore, the user can practically organize them (for example, according to their topics) taking advantage of the same ontology. In order to improve the basic clustering techniques (which cluster only using document/term matrices and lose much of the implicit information contained in the language) some background (domain) knowledge is introduced in the process. Doing so, it is possible to provide "subjective" views onto document collections. For example, highlight differences and similarities on the content, its presentation form, or on the skills needed to approach it.

e) An ontology evolution component discovers changes and trends within the field of interest thanks to ontology learning methods. For example, it supports the introduction of new concepts and checks for inconsistencies, it recognizes "concept drifts" (that is, the change of meaning of concepts in constant flux, such as ‘Semantic Web’), it supports versioning ontologies (also if they change, it can still be useful to relate them), and, in general, it accompanies the user in the ontology lifecycle.

In conclusion, the Courseware Watchdog integrates a complex series of functionalities that represent quite well the whole authoring process a teacher is normally involved with.

Also Aroyo and others (Aroyo and Mizoguchi, 2003,Aroyo et al., 2003) have worked in the same direction, attempting to fulfill the “constant requirements for educational content flexibility and adaptability, for learning objects and structures reusability and sharing”. At the basis of their approach there is the awareness that ontologies are the ideal infrastructure for integrating intelligent systems and enabling knowledge sharing.

In this context, they propose an *authoring task ontology (ATO)* as the main solution for supporting the authoring process in all its activities, and for providing it with a methodology and with a vocabulary. The ATO is compounded of *authoring activities*, *sub-activities*, *goals* and *stages*, within a framework that formalizes the semantics of the whole authoring process. The authors have instantiated this ontology-based approach to courseware authoring by ‘upgrading’ two existing web-based systems, SmartTrainer (built at Osaka University, Japan) and AIMS (realized at University of Twente, The Netherlands). They show how the whole authoring process can be organized into a three layered framework (see fig. 2-13). At the top level they pose the definition of the so-called ‘static knowledge’, that is, the curriculum organization with instructional design models. The middle level corresponds to the ‘dynamic knowledge’ definition, that is, the tutoring strategy adopted in order to tailor the learning to the learner. Finally, at the bottom level there are the specific instructional systems instantiating these ideas, SmartTrainer and AIMS .

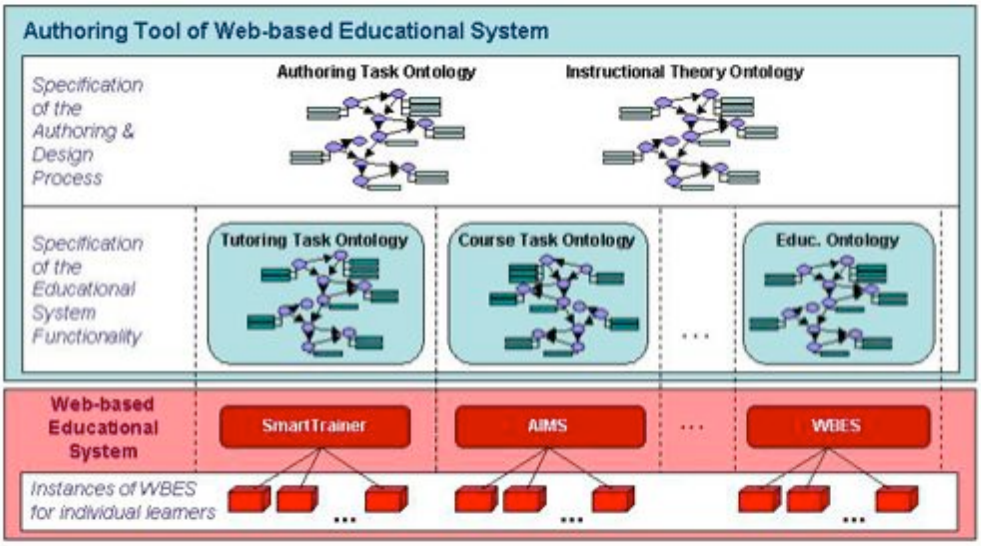

Figure 2-13 – Ontologies involved in the authoring process (from Aroyo, 2003)

The authoring task ontology, it is worth noting, remains the main focus of the authors. In fact, by modularizing and specifying all the authoring activities at the maximum level, the system’s domain, the educational strategy and the educational goals could remain totally independent.

2.4 Navigating through semantic spaces: research directions

Similarly to the systems described in the previous section, many of the research applications we are about to discuss in the next paragraphs were created within an educational context. Nevertheless, we decided to put them into a separate section, because we want to focus our attention on another important feature they exhibit.

That is, all these systems employ semantic technologies in order to *improve or facilitate navigation and presentation of digital resources repositories*. This is achieved in various ways and at different levels; but to some extent, we believe it is possible to look at all these systems as attempts to transpose (more or less explicitly, and more or less faithfully) the *digital narratives* approach (see section 2.2.5) into the Semantic Web world.

In the order presented, we are going to focus our attention on applications for semantic browsing (section 2.4.1), faceted browsing (section 2.4.2), hypermedia discourse-generation (section 2.4.3) and ontology-based hyperlinking (section 2.4.4).

2.4.1 Semantic browsing

The underlying research question driving the development of these systems can be expressed as follows: if resources can be adequately annotated and retrieved, *the problem becomes how to browse them in an efficient and focused manner*.

As we have seen, digital narratives (see section 2.2.5) address this same problem, by creating mechanisms that allow diverse browsing possibilities (i.e.

'discourse' structures) over the same set of atomic resources (i.e. the 'story' items).

It is worth remembering that also a number of systems in the hypermedia area have addressed the same issue, mainly in order to achieve an increased independency between the navigational level and the resource level. For example, the authors of Walden's Paths (Dave et al., 2003,Shipman III et al., 1998) argue for the supremacy of *path-centric browsing* as an independent browsing mechanism capable of overcoming many classic navigation-related problems, such as 'cognitive overload' and 'disorientation' (Conklin, 1987). According to the authors, *paths* are defined when people "select a subset of the information contained in the hypertextual networks and organize it for use in specific contexts of interest to them and their readers". Obviously, this type of solution is strictly dependent on a manual creation of each single path available. Also, the information implicitly contained in the sequencing of the elements is not accessible as such, therefore, it is hardly re-usable.

In more recent times, *conceptual navigation, supported by ontological engineering*, has emerged as a solution to this problem (Crampes and Ranwez, 2000). Thanks to *domain* ontologies and *argumentative* ones, the links between resources and their narrative or pedagogical roles are computed from their description in a formal conceptual language, and may vary according to the situation. The navigation, therefore, happens at the conceptual level and can provide features such as *conceptual expansion* (some sort of lateral browsing, that takes into account concepts not directly related) *forward conceptual navigation* (a process similar to the free navigation in a hypertext, with the difference that is based on concepts) or *conceptual specification* (the retrieval of the direct sub-concepts of the initial one).

In all these cases, the navigation is called *ontology-supported*, since "the ontology of the domain has been used only for indexing in order to be sure that the resources and the users' objectives are described with the same concepts and relations" (Crampes and Ranwez, 2000).

Instead, when the paths the user can choose from are partially determined by the system, in a narrative or pedagogical ontology, the navigation becomes *ontology-driven.* As the authors claim, their "central idea for ontology-driven conceptual navigation is to design an architecture where the engine only relies on ontologies for selecting resources, ordering them and adding a narrative/pedagogic intention during the linking process" (Crampes and Ranwez, 2000).

Systems of the first kind ('ontology-supported'), can be the one presented by Alfaro and colleagues (Alfaro et al., 2003), which supports the browsing of multimedia resources at the conceptual level. Other examples include Poncelet (Habel and Magnan, 2007), a system that lets users learn about mathematics by navigating an ontology-based multimedia environment; or the work of Becker and colleagues (Becker et al., 2003), an analogous system which is instead tailored to enhance the navigation of intra-organizational information systems.

We are more interested, instead, on systems of the second kind ('ontology-driven'), since they provide more elaborated narrative facilities which can resemble the strategies adopted by the digital narrative systems introduced in section 2.2.5.

Celino and Della Valle (Celino and Della Valle, 2005) present an ontology-driven environment inspired by the metaphor of the *journey*. The authors call a ‘hyper’ environment “any information system in which resources are described in a machine-processable way”, and subsequently introduce the notion of a *vehicle* as the “necessary tool to navigate effortlessly across a hyper-environment and to follow the most opportune path to reach the needed information”. Essentially, *vehicles* appear to be the chosen metaphor abstracting possible ways to browse a semantic repository.

These ideas are then instantiated with a framework that supports the building of Semantic Organizational Information Portals (SOIP-F). Such a framework, claim the authors, separates the modeling of *domain* information, *navigation*, *access* and *presentation* knowledge. This is achieved by means of a *portal ontology* that includes portal-dependent terminology: structural terms such as *entity* or *component*, navigation terms such as *contains* or *related_to*, access terms such as *next* or *down* and presentation terms like *title*, *text-box* or *image*.

In a real-world scenario, the framework is used in conjunction with some *domain* knowledge too (i.e. an ontology). The interdependence among the various representation layers is thus achieved by “mapping the domain terminology into the portal terminology, creating, in a bottom-up approach, a relation between domain-dependent terms and portal-dependent terminology”. In conclusion, as a result of this mapping process, the authors can define different *vehicles* for browsing the information space.

Another interesting system implementing ontology-driven navigations is Story Fountain (Mulholland et al., 2004,Mulholland et al., 2003). This is a tool developed to *support a community in the exploration of digital resources,* specifically stories. The background approach is constructivist, in as much as it lets the users engage with the subject matter, make their own interpretations and basically learn through a story sharing process. Users, in fact, 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.

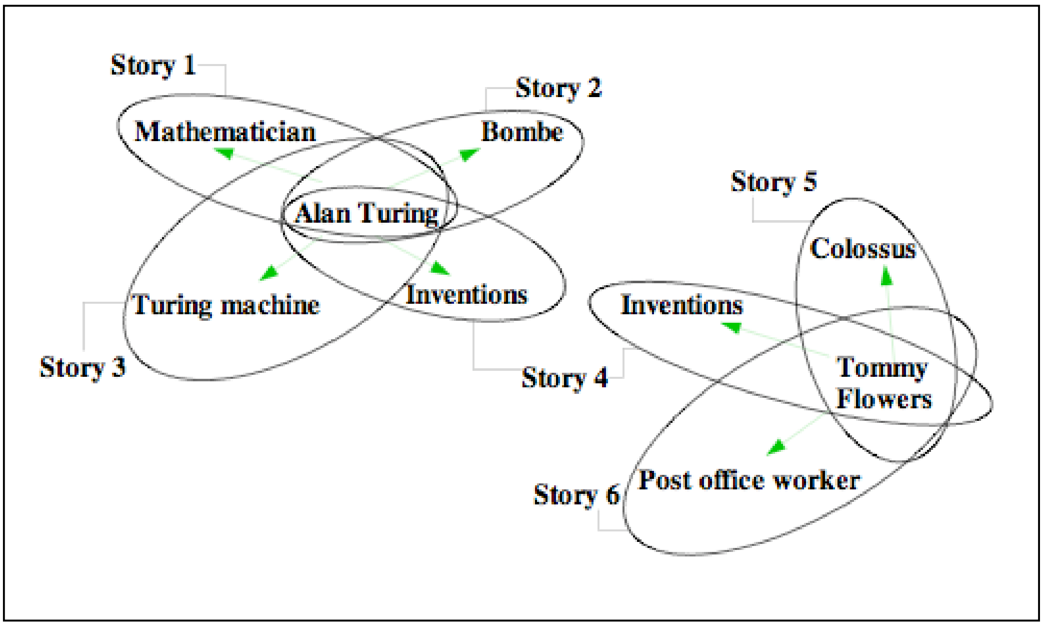

Fig. 2-14 - Example of a conceptual path in Story Fountain (from Mulholland 2004)

Thanks to a *domain* ontology and a *narrative* ontology, the different stories annotated and stored in a database are later recollected in an intelligent way. Compared to a simple string matching retrieval, Story Fountain provides a great improvement towards the understanding of the stories; in fact, it generates semantic navigational paths as a result of the novel connection of different concepts.

Of great interests are the six *exploration facilities* (corresponding to concept-pathways) that the system provides, and that are based on the narrative ontology:

1. *Story Understanding*: a view that highlights the conceptual structure of a story in terms of its central characters, events, physical objects and themes (these are classes in the story and narrative model).
2. *Concept Understanding*: a facility that collects all stories containing a selected concept, in order to do a comparative study between stories (see fig. 2-14).
3. *Concept Comparison*: it selects stories related to multiple concepts in order to do a comparison between concepts.
4. *Concept Connection*: it automatically draws pathways between stories thanks to the interrelations of concepts and events defined in the ontology.
5. *Story Mapping*: it gives a story perspective depending on a selected concept.
6. *Event Mapping*: similar to Story Mapping but with the properties of events instead of stories.

At a higher level, the kind of interaction observed in the users while employing these facilities are classified by the authors into four exploration processes: *accumulation* (the aggregation of information of a particular type across different materials), *association* (identification of contingencies between different concepts or events), *induction* (usage of source materials to support hypothesis) and *information gathering* (basic exploration process without a detailed aim).

In conclusion, although 'ontology-supported' navigation mechanisms provide a rather basic mechanism for browsing a semantic space (i.e. by following precisely the taxonomic structure of the underlying ontological 'skeleton'), this is not the case for 'ontology-driven' navigation mechanisms.

In fact, systems like the just mentioned Story Fountain can provide very complex and domain-specific *ways to traverse* an information space. In particular, these mechanisms can be tailored to specific users' needs (e.g., learners) and domain characteristics (e.g., academic disciplines).

We will return on these topics later, given that our 'learning narratives' approach can also be classified as a type of 'ontology-driven' navigation.

2.4.2 Faceted browsing

A more general-purpose approach to the problem of navigating semantic spaces is *faceted browsing* (Yee et al., 2003). This is a technique for exploring structured data-sets based on *facet theory* (Ranganathan, 1990). As Oren and colleagues explain (Oren et al., 2006):

> "In faceted browsing the information space is partitioned using orthogonal conceptual dimensions of the data. These dimensions are called facets and represent important characteristics of the information elements. Each facet has multiple restriction values and the user selects a restriction value to constrain relevant items in the information space. The facet theory can be directly mapped

> to navigation in semi-structured RDF data: information elements are RDF subjects facets are RDF predicates and restriction-values are RDF objects."

In the case of semantically annotated repositories, faceted browsing has recently emerged as an alternative to semantic search (Guha et al., 2003). At the roots of this adoption stands a widely accepted view which classifies the possible information needs of a user (Marchionini, 2006) (C.W. Choo et al., 2000) (Wilson, 1994). Basically, this theory tells us that it is possible to draw a distinction between 'lookup' queries, where users are looking for precise answers to a well-defined questions, and more 'explorational' queries, whose objectives include also other more complicated activities such as *comparison*, *analysis*, *interpretation* and *discovery* of new information.

While in the case of 'lookup' queries classic string-search engines may perform well, in the latter case it is more profitable to employ techniques which retrieve a set of possibly relevant documents and then let users interact with the results set so to refine and better understand their information need. This new paradigm has been called *exploratory search* (Marchionini, 2006).

In this context, faceted browsing has been used by several authors as the main technique for supporting this type of exploratory tasks. In particular, faceted browsing proved to be useful when dealing with semantic repositories whose precise structure is unknown to the user. We will now examine a few systems which showcase such techniques, highlighting their most salient features.

/facet (see fig. 2-15) is a tool developed in the context of the dutch MultimediaN E-Culture project (Hildebrand et al., 2006,Schreiber et al., 2006) for supporting the browsing of large virtual collections of cultural-heritage resources.

By using an architecture fully based on open standards (XML, RDF/OWL, SPARQL) /facet allows the navigation of any user-defined semantic repository. More precisely, this tool aims at supporting both the a) annotation of web resources present in the dataset (in the available demo, images) and the b) search and presentation/visualization of the same resources. Among its principal features, the following ones should be mentioned:

- It allows the *navigation of several repositories* (i.e. different knowledge models) at the same time, without needing any previous manual alignment of them. This is done by regarding the *rdf:type* property just as another facet, which users can select to narrow down the search space.
- It provides a *keyword search functionality* so to facilitate users in finding facets and items of interest; the authors of /facet describe such keyword search as ‘semantic’ fundamentally because its results are presented in a tree structure, thus highlighting what position they have in the original rdf model. The keyword search can be used for investigating both the *types* of facets available (e.g., searching for all facets containing the word “style”) and the *items* within a specific facet (e.g., searching for all the paintings containing the word “Arles” within the ‘post-impressionist style’ facet).
- It allows users to *query over multiple resource types*, by searching for common properties of the different types at run time.

- It supports the creation of *configuration files* (in RDF), where additional relationships among facets and types can be defined; this feature aims at letting users personalize the hierarchies shown in the interface.
- It provides facet-specific *interface extensions*: e.g., a timeline view for historical data, or a map-based visualization for geographical ones.

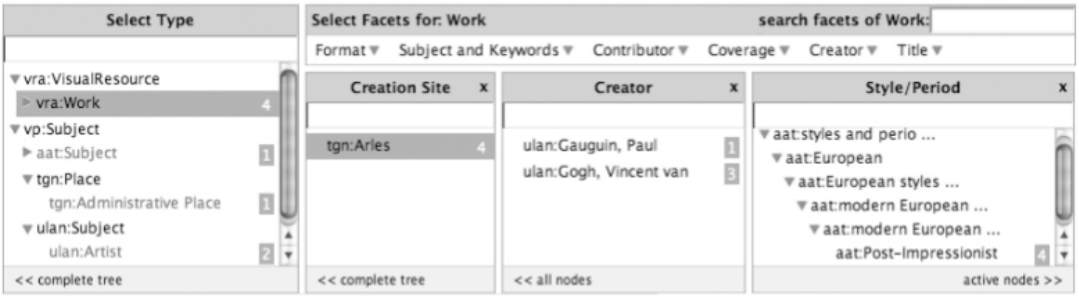

Fig. 2-15 - Faceted browsing in /facet (from Hildebrand 2006)

The CultureSampo finnish cultural heritage portal (Eero Hyvönen et al., 2007,Eetu Mäkelä et al., 2007) also adopts faceted browsing for letting museum visitors interact with a large collection of semantically annotated items. Their approach is slightly different from /facet, as far as they intend to shift the focus from object location to the automatic creation of *domain-centric presentations*. It is worth quoting the original passage where this approach is stated:

> “Our idea is to let users create virtual exhibitions that mimic the way real museums are organized, containing themed exhibition rooms of items and displays that together, through the objects, tell the story of a particular subject”.

The CultureSampo portal is aimed at the general public: in order to facilitate the task of starting an exploratory search, users can specify their interests by operating on a structured query-composition interface (e.g., "Tell me about *item type* related by *role* to *domain concept* organized by *classification+role*"). The results thus obtained are then visualized using a faceted browsing approach.

In particular, the authors highlight a difficulty arising from the specific domain they are dealing with: cultural domains typically have many content *types* with even more *properties* attached to them. As a consequence, the final visualization can be very complex because the combinatorial space is too vast to be represented effectively. As a solution, they propose a 'domain-centric' approach to faceted browsing. In this case, the "properties are relegated to a secondary role, and the views were built instead based on the ontological ranges of those properties, i.e. the set of topical domain ontologies".

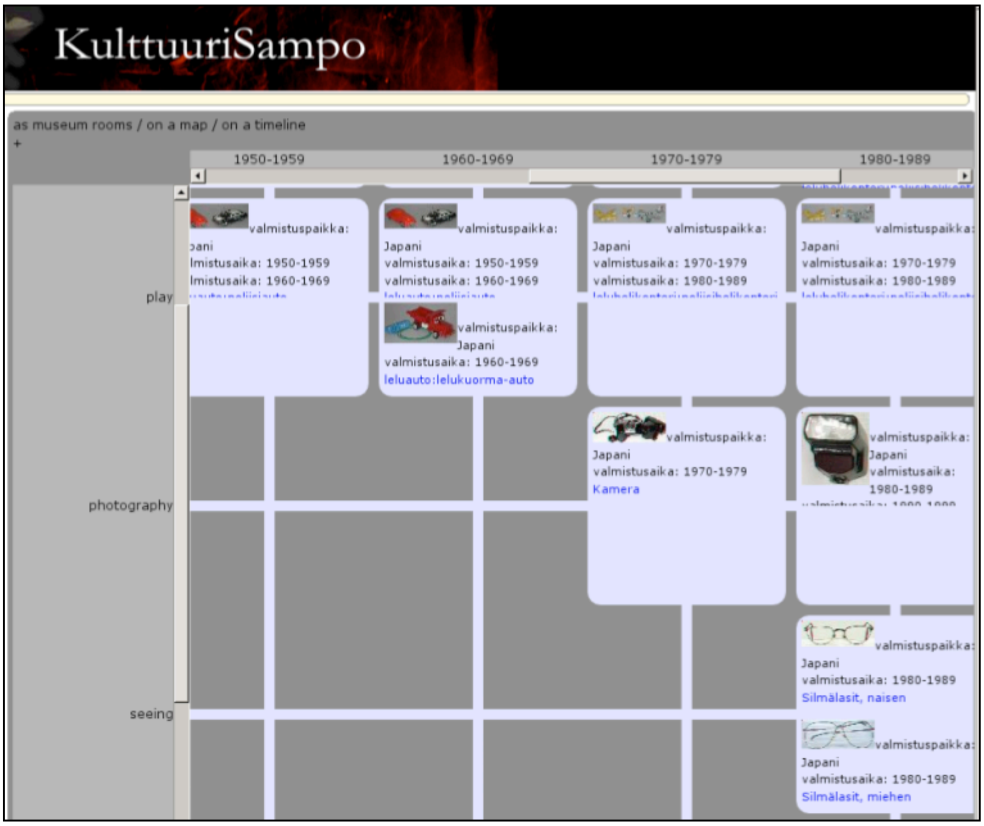

Fig. 2-16 - The virtual museum rooms in CultureSampo (from Hyvönen 2007)

Also, it is worth noting the way the visualization attempts to reproduce a typical museum's structure, with themed floors and rooms of exhibits. Moreover, within a room's visualization, the results are presented in a two dimensional matrix whose row and columns "are comprised of a flattened list of concepts in the two domain facets chosen for organization" (see fig. 2-16). By doing so, each cell in the matrix corresponds to rooms combining two themes (e.g., "18th century agriculture" or "18th century industry").

Exhibit (Huynh et al., 2007) has quite a different approach. In fact, the faceted browsing functionality it offers is not the main purpose of the system. The authors describe it instead as a "publishing framework designed to do for structured content what HTML has done for unstructured content: lowering the barrier to publishing while offering a high level of control over presentation".

The target audience is also quite different: mainly, non-expert individual users with moderate quantities of semi-structured information (an 'exhibit' can handle up to a few hundreds items, say the authors) who search for an easy way to publish their data.

By doing so, according to Huynh and colleagues Exhibit accomplishes the double purpose of a) facilitate users in making their semi-structured data available and b) foster the development of the Semantic Web, because such an easy-to-use publishing framework is likely to augment the number of 'average' web-pages' creators who want to share their data-sets. .

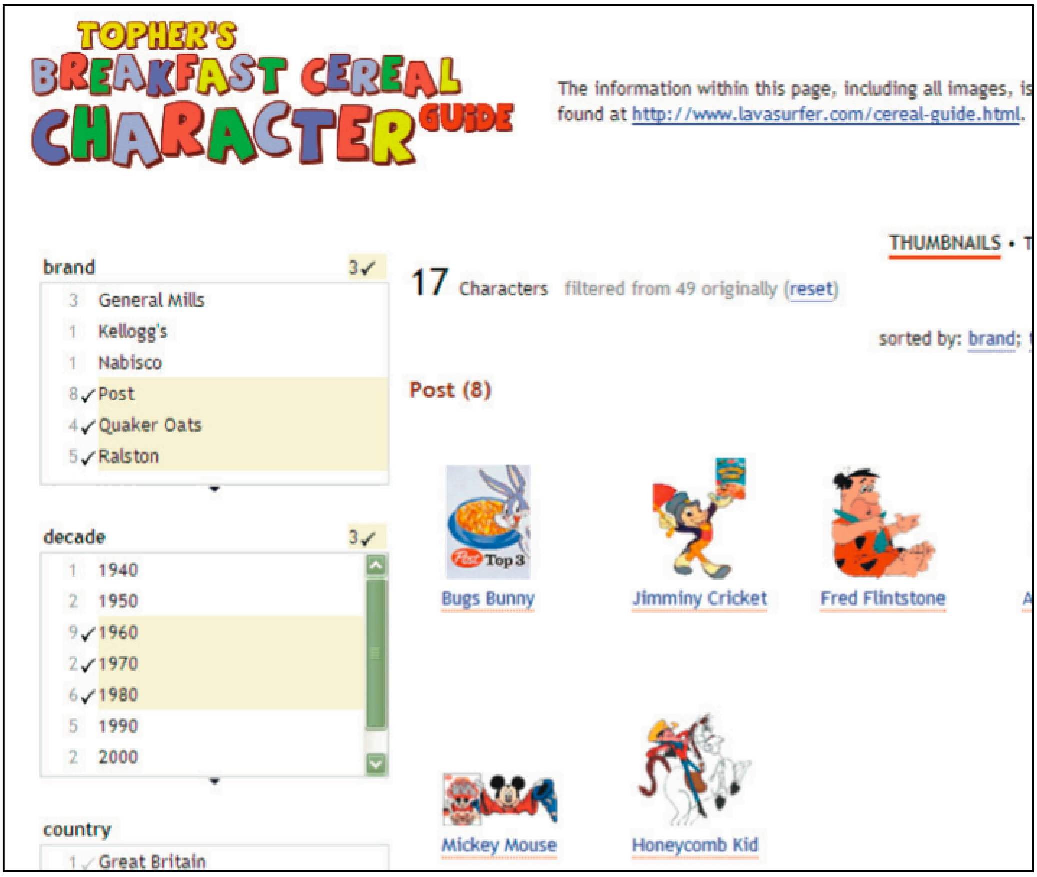

Fig. 2-17 - An example of exhibit's faceted browsing interface (from Huynh 2007)

Among the other features of Exhibit we should highlight, there are the following ones:

- It is conceived as a *lightweight* application: there is no need to install anything, configure or maintain the source code. Authors can tweak the system, but even without doing so they can still create an exhibit out of some semi-structured data;
- The faceted browsing interface (see fig. 2-17) is customizable, and it offers mechanisms for personalizing how to render single items ('lenses') or a set of items ('views');

- As for the data model, Exhibit can read data in its own JSON (JSON.org, 2008) format. However, translators to/from RDF and other standards are available;
- All the data shown in an exhibit can be *exported* by clicking on the ‘copy’ button: this generates a bunch of code (RDF, JSON etc.) describing the items selected, thus facilitating reuse in other applications.

mSpace (Schraefel et al., 2003,Schraefel et al., 2006) is defined as an “interaction model and software framework that brings together a variety of mechanisms to improve access to information by supporting multiple ways of exploring the information itself”.

Similarly to the previous applications, mSpace aims at supporting exploratory searches in unfamiliar domains codified according to SW standards. This is done by means of a highly interactive interface where users can ‘play’ with domain objects’ metadata, thus triggering different visualizations of them.

A first instantiation of the mSpace framework was done with the Classical Music Explorer (Schraefel et al., 2005) (see fig. 2-18).

Fig. 2-18 - Faceted browsing with mSpace (from Schraefel 2005)

This application lets users learn about classical music's composers, songs, styles etc. using a sophisticated faceted browsing interface. It is worth highlighting the innovative features it showcases:

- *Audio preview cues*: e.g., by clicking on a song's preview cue users can quickly sample parts of the domain and discover what are the areas of interest. Notice that preview cues, differently from common music-store softwares, are available not only for 'song' objects but also for 'styles', 'authors', 'arrangements' etc. In such cases, users are given various choices about which 'cue' to play.
- *Spatial layout*: in a spatial multicolumn layout the results of selections made in one part of the display are displayed in another part of the same display, not on a different page. The main advantage of this technique is that users are helped in maintaining awareness of *contextual* information.

- *Slices*: in order to support users' interaction with the multiple dimensions an information space (in this case, classical music) may have, mSpace uses the 'slice' metaphor. That is, every dimension (i.e. every facet) is represented through a slice. Slices are "arranged from left to right, in columns, creating a hierarchy where the left-most column is the top level of the hierarchy and the right-most is at the bottom". An important feature of slices is that they are dynamic: users can alter them by rearranging, adding or subtracting dimensions. This is in fact the main mechanism by which they can explore the domain and organize it according to their interest.
- *Favorites*: users can add instances of interests to a 'favorites' area. Also, once they later retrieve the saved items, by clicking on it the path taken (i.e. dimensions selected) to record that item is made available. This feature, claim the authors, let users focus on the domain exploration, rather than having to remember what they wanted to look at further on.

Finally, the Collex system (Nowviskie, 2005), developed in the context of the NINES project (Nowviskie and McGann, 2005) (a major collaboration among humanities' institution for sharing digital resources and creating metadata about them) is targeted to humanities scholars who increasingly need to perform the following operations online:

1. Collect, tag, analyze and annotate trusted objects (digital texts and images vetted for scholarly integrity).
2. Reorganize and publish the same objects in fresh critical perspectives.

3. Share these new collections with students and colleagues, in a variety of output formats.
4. Without special technical training, produce interlinked online and print 'exhibits' using a set of professional design templates.

Collex (see fig. 2-19) was created as the answer to such needs. Essentially, it is a faceted browsing application based on Semantic Web standards which is very similar to /facet. A quite interesting feature is the possibility of creating 'collections' of objects of interest and then share them (also as RDF representations) with other people.

In general, when compared to the systems presented above, Collex does not present any major technical improvement. Nonetheless, we decided to include it in our review for the simple reason that this is the first serious attempt to 'export' semantic navigation technologies to the humanities' area. As we will see, this same intent is characterizing our work; as we are going to focus our attention to the development of an *ontological model* usable for navigating semantic repositories in the humanities domain (in particular, philosophy).

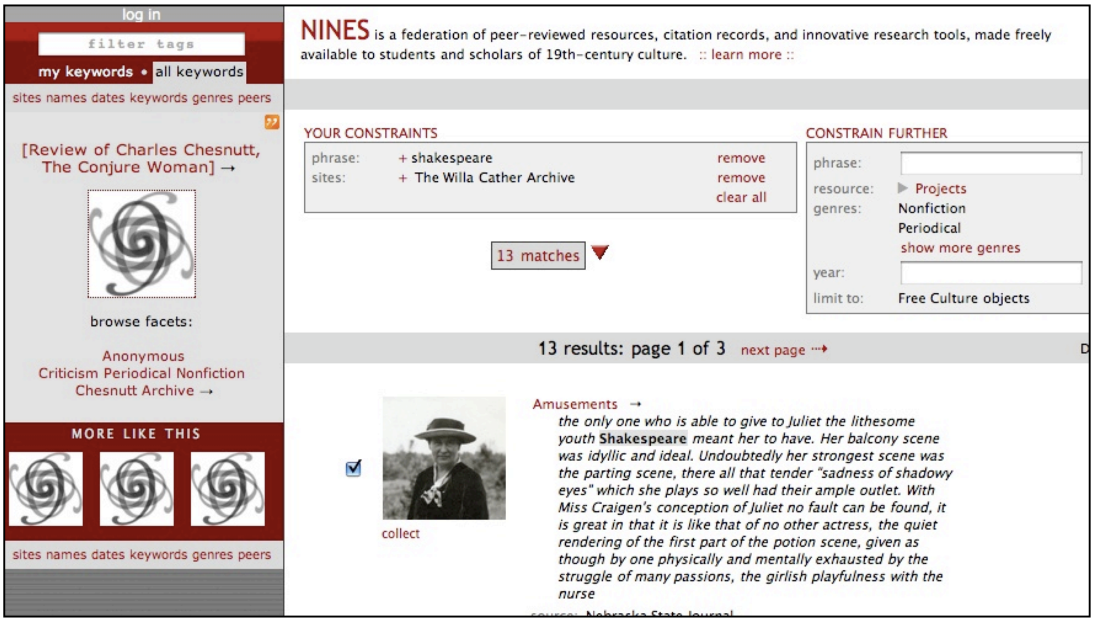

Fig. 2-19 - The Collex faceted browser for humanitites' resources (from Nowviskie 2005)

In conclusion, faceted browsing is a very powerful technique for navigating large semantic digital repositories. In particular, as stated by various researchers, this technique is best suited for the *exploration* of unknown knowledge domains. In fact, by means of highly interactive visualizations' 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 system, in our opinion, is linked to their very best feature. That is, being largely non-domain specific and allowing navigations 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. Just to mention an example, in the field of philosophy we might want to present resources in a way that highlights the 'intellectual lineage' of an idea, or the 'historical attempts to solve a problem', without having this information readily at hand in the form of a facet.

In other words, faceted browsing is not suited for supporting situations where we need to perform some *reasoning tasks* on the semantic layer describing our resources. In such cases, the approach discussed in section 2.4.1 ('semantic browsing') is definitely more successful. In the long run, a combination of the two seems also a very interesting path to take.

2.4.3 Hypermedia discourse generation

We will now look at a class of ontology-based applications where the accent is posed not only on the *navigation* facilities but also (and sometimes, even primarily) on the *discourse generation* ones. Such applications make use of *domain* and *narrative* knowledge, codified in different ontologies, with the purpose of organizing a set of raw materials into a coherent presentation. We will see in more details how this happen by analyzing four exemplar applications.

The Topia project (Rutledge et al., 2003) has been developed and used in the context of the Rijksmuseum Amsterdam, to support visitors in the browsing process of the works of Van Gogh. Topia is focused on the production of *hypermedia presentations from the semantics of potentially unfamiliar domains* (in their example, expressed in RDF). Although the authors recognize the necessity of human insight in order to generate a story, they still feel that there is a "subset of narrative and discourse concepts that one can automatically derive from semantics", and that includes, for example, the *order* of a presentation and the *grouping* of components into sections.

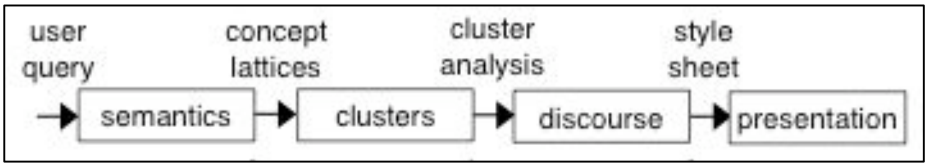

Figure 2-20 – The four phases in the Topia system (from Rutledge, 2003)

This approach, therefore, being strongly domain-independent and computable, relies on the *clustering* of similar concepts and on their weight-assignment based on simple features, like cluster size. In order to achieve this, the authors highlight four different phases (see fig. 2-20).

After the user prompts a query (phase 1 – semantic processing), the system tries to match it with the items in the RDF repository, and returns them together with the property assignment every item has (since RDF, as explained above, stores knowledge in the form identifier/property-type/property-value).

This set of collected items is then passed to the clustering algorithm (phase 2 – concept lattices), that has the function of looking for patterns in the graph (the RDF bundle of items) that act as landmarks for important locations. A *cluster* is a node with close proximity to a relatively large number of the originally selected nodes in the graph. A *concept lattice* is a particular clustering technique, and the authors use it in order to group the retrieved results into a partially ordered set of item-set/property-set pairs. Within this set, a *concept* is defined as the union of the items that share the same property, and it constitutes the basis for an informative structure around the items.

This totally “syntactic” process (since it does not make use of any explicit domain knowledge) generates groups of items that are weighted depending on their proximity to the cluster’s centre, and are passed to the next module in order to generate a *structured progression* (phase 3 – clustering analysis). In addition, the system allows the user to specify the significance of certain concept’s properties, thus modifying their weight in the final presentation. At this point, clusters have some measured rating of importance, which determines what type of discourse construct each cluster becomes and what order components are presented in. The components of structured progressions (the “subset of narrative and discourse concepts” we were talking about at the beginning) are the following, in order of importance:

I. *Hierarchical Structure*, similar to the division into sections and sub-sections in a textbook. Only clusters that are significant enough become part of the discourse hierarchy.

II. *Meaningful Order*, that is, the sorting of groups based on the minimum, maximum or average of their items’ value for a sorting property (that can be entered by the user).

III. *Recurring Themes*, namely properties shared by multiple items distributed through the discourse hierarchy.

IV. *Tangents*, the remaining least significant clusters that are not important enough to appear in the primary flow of the resulting presentation.

Finally, the structured progression has to be converted into a hypermedia presentation, since otherwise it would remain just an abstraction of how a presentation should be, without the details of its implementation (phase 4 – style sheet presentation). In order to do so, claim the authors, it is useful to follow a principle of *discourse perceptualization.* Such principle states that the user should perceive, at every point of the presentation, the overall structured progression and the context of the current point within it.

Geurts and others (Geurts et al., 2003), instead, describe a system that makes use of different ontologies in order to generate a complex presentation design. The background assumption of their work is that since also *information presentation design* is an inherently knowledge-driven process, it can partially be automated and benefit from the exploitation of SW technologies. In their system the user does not only enter a query, but also specifies some characteristics of the final presentations, like the *genre* (biography or CV) and the *medium* (printed paper, hypermedia presentation). This is possible because the system takes advantage of various ontological representations: first of all of the *domain*, then of *discourse* and *narrative*, and finally of *design* and *multimedia* knowledge (see fig. 2-21).

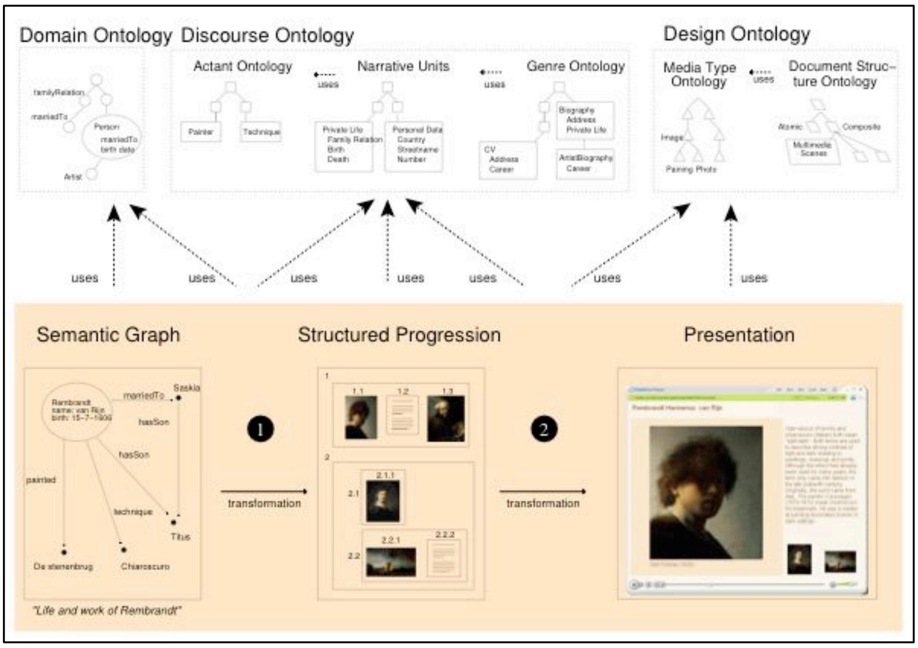

Figure 2-21 – The two-phase presentation process (from Geurts, 2003)

So, if multimedia items are properly annotated, they can be matched by a user-query and retrieved in the form of a semantic graph (e.g., in RDF). At this point, *both* the domain ontology and the discourse and narrative one are used to deploy a *structured progression* (analogous to the one presented in the previous system's analysis). However in this case, the clustering phase is not necessary anymore since the useful knowledge is already codified and available, thanks to the domain ontology. This allows the development of patterns within the retrieved RDF's bundles. For example, within a domain of painters and works, similar properties of the works can lead to novel connections between different authors, and so on. The domain ontology specifies the structure and the important relations in the domain, while the discourse one provides the rules to reason on the domain information and connect far-away resources into a coherent narrative.

In a second phase, instead, the structured progression is transformed into a final multimedia presentation. This is achieved in two steps: first of all a document structure is created, where all the decisions about the output medium are made explicit. Secondly, when all the detailed layout and formatting features are made clear, the document structure is transformed into a tree of formatting objects.

A slightly different system is Artequakt (Alani et al., 2003), developed at the University of Southampton, since it *combines a powerful information extraction tool*, to populate a knowledge base with information gathered on the Web, *with a presentation module based on templates*, to format the final output for the user. In this case, as we can see from figure 2-22, there is only one ontology, the domain one, which is used to provide directions for the information extraction phase, and structure for the organization of the retrieved material.

Artequakt seeks for information about artists on the Web, stores it and re-assembles it to generate personalized narrative biographies. The presentation phase basically rests on human authored biography templates (authored in the Fundamental Open Hypermedia Model -FOHM- developed by the same university), where the basic structure used is a *Sequence*, that is, a list of queries to the knowledge base that have to be instantiated and inserted into the biography in order. In addition to this, templates may include their own text (like a sub-heading title), and some patterns to construct basic sentences if only the data are available.

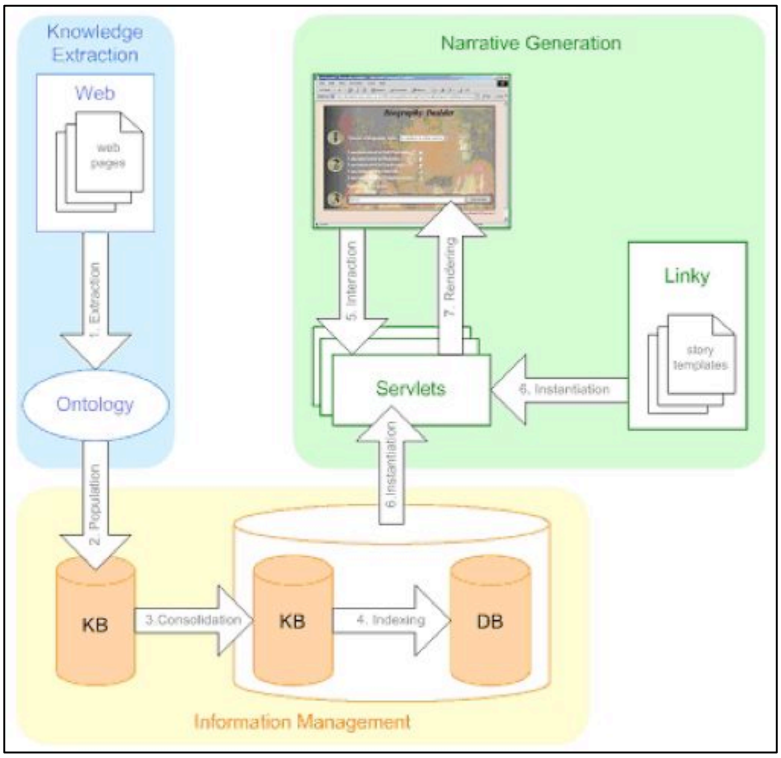

Figure 2-22 – Knowledge extraction and presentation in Artequakt (from Alani, 2003)

Obviously, from the narrative point of view (cf. section 2.2.4), the ontological organization of the information corresponds to the *story* level, while the *discourse* level corresponds to the fixed templates provided by the authors, thus not much reasoning or intelligence is provided.

Another system that *derives semantics from shallow annotations and then presents it in a personalized manner* to the user, is the one introduced by Little and colleagues (Little et al., 2002) . In this case, the annotated data are retrieved from the Open Archive Initiative (OAI), a community that has defined an interoperability framework to facilitate the sharing of metadata, expressed in the Dublin Core (DCMI, 2008) format.

Since from a basic keyword search on the metadata of this library too many items are retrieved (and between them, often there is no relation at all), the strategy adopted is to let the user *direct* the search process. For example, the first skimming of a large number of items is performed when the user selects only the relevant results, and the search process is iterated. During this operation it is possible to *infer* semantic relationships between resources and step-by-step *focus* exactly on the items the user is looking for.

The semantic relationships are derived directly from the metadata schemas (for example, the system may take a *dc.contributor* value for a chosen resource and search for resources which have *dc.subject* equivalent to this value), consequently, they will suffer from two internal limitations of this schema. First of all, since they are very specific to the Dublin Core metadata associated to the set of acquired media objects, they are hard-wired into the system as a set of pre-defined rules (and cannot be encoded in an ontology); secondly, they rely on the poor expressiveness of DC metadata.

An ontology is used, instead, in order to provide knowledge relevant to the <u>presentational</u> phase: the task to complete is to map semantic relations to spatial/temporal relations. The most relevant constraint, here, is the fact that while the possible semantic relations in a domain are infinite, the number of possible spatial and temporal relationships is limited. Therefore, the ontology is derived from the reduced set of top-level MPEG-7 (Multimedia Description Schemes specification) semantic relationships. For example, “X describes Y” (semantic level) is mapped to “X annotates Y” (MPEG7 specs) and then to “spatialBelow (X, Y), spatialAlign (X, Y)” (temporal/spatial level).

Finally, the Cuypers Presentation Generator (Ossenbruggen et al., 2001), a five-layers technology to translate the semantics of a presentation to its final output, is employed to send the results to the user's Web browser or media player.

2.4.4 Semantic hyper-linking

We conclude this section about systems for 'semantic navigation' by mentioning a group of research efforts achieving this type of navigation facilities in the form of 'semantic hyper-linking'. This class of systems have their main characteristic in the usage of ontologies so to *enrich already-existing web documents with additional relevant links*.

In general, the way this is achieved is by operating a two-phase process: firstly, scanning the original document for potential link-anchors which match the ontological knowledge; secondly, constructing links to resources which have been previously attached (manually or automatically) to the ontological descriptions. We will analyze two examples of such systems.

Cohse (Carr et al., 2001) was a seminal work in this direction. As the authors claim, *the union of an ontological reasoning service and an hypermedia link service generates a 'conceptual hypermedia' system*. This type of technology is expected to "enable documents to be linked via metadata describing their contents and hence to improve the consistency and breadth of linking of WWW documents at retrieval time (as readers browse the documents) and authoring time (as authors create the documents)". The architecture of a conceptual hypermedia system can be broken down in the following key-components (see fig. 2-23):

a) The *ontology* service manages the ontologies and maps between natural language terms and a concept graph.

b) The *resource* service obtains web pages related to the ontological concepts.

c) The *link* generator service is at the centre of the system: it looks for ontological terms in the document (by contacting the ontology service); it queries the resource service for link destinations relevant to the selected ontological concepts; it creates the 'enriched' document using further 'editorial constraints' and also decorating it with several metadata (i.e. language terms from a specific ontology).

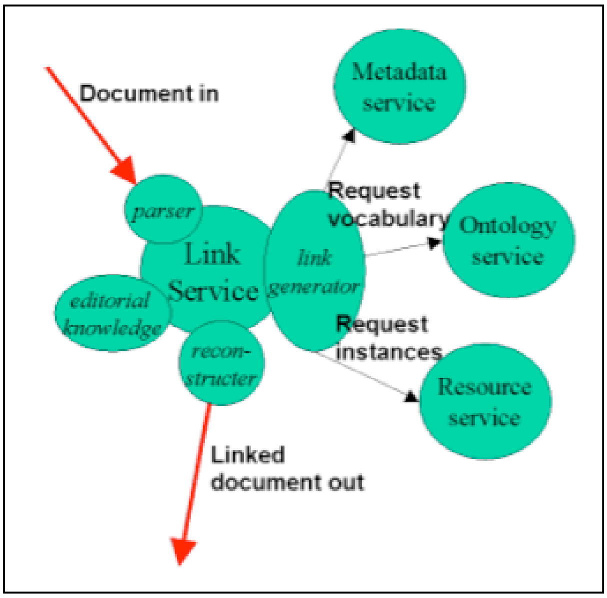

Fig. 2-23 - Schema of the architecture in Cohse (from Carr 2001)

As an example of an instantiation of this framework we shall mention GOHSE (Bechhofer et al., 2006). Here the domain knowledge is given by a Gene Ontology, which is used to enrich biological documents with dynamic links to relevant literature.

From the architectural point of view, it is important to remember that COHSE operates as a server side application: by means of a *proxy*, the pages users want to be enriched are first retrieved, then augmented with new hyperlinked and passed back to the client.

A slightly different approach to semantic hyper-linking is the one proposed with Magpie (Domingue et al., 2004,Dzbor et al., 2003) or VIeWs (Buitelaar and Eigner, 2005).

The main difference relies in the fact that Magpie (and similarly, VIeWs) is a browser plug-in which lets users select one or more 'ontological spectacles' to be used with the active web-page. That is, users can select one or more ontologies so to highlight concepts belonging to a knowledge domain on the webpage they are reading.

By switching between one 'spectacle' and the other, users can benefit from the association of different 'semantic layers' to the same web resource. In practical terms, the enriched documents are not created every time through a call to a proxy service (such as in Cohse), but are instead created in the browser at *run-time*. Accordingly, as the authors claim, "Magpie is essentially a bridge, a mediator between formal descriptions used by the ontology-based service providers and semantically unstructured web documents".

Another important feature to highlight is the fact that once the semantic entities on a web page are annotated, the contextual (right-click) menu of a web browser is overridden by an *on-demand services menu* (see fig. 2-24). This menu is context-dependent, as could be expected, but with the particularity that in this case the context is *semantic* i.e. it is defined by the membership of a particular entity to a particular ontological class.

This feature is quite interesting, as it embodies the idea of a semantic navigation through context-dependent 'pathways' mechanisms.

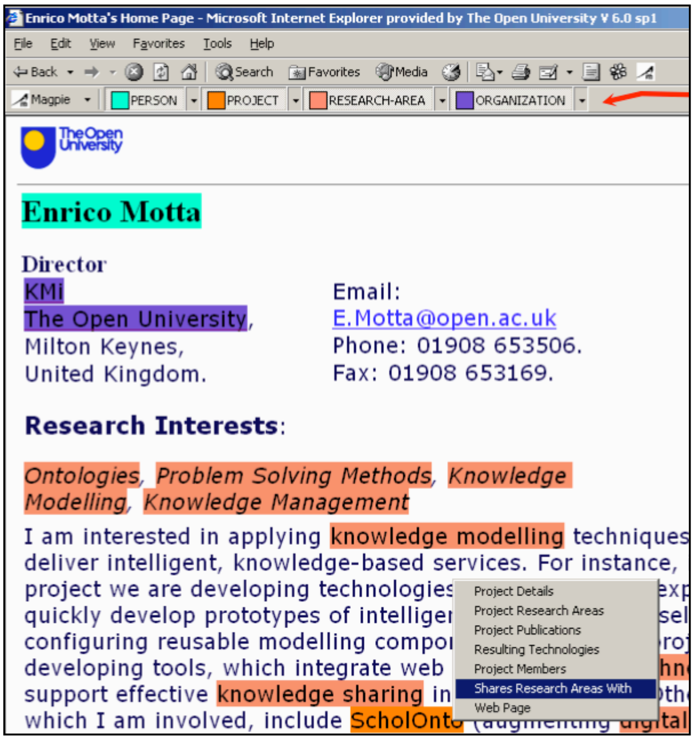

Fig. 2-24 - Semantic navigation in Magpie (from Domingue 2004)

2.5 Summary and gap analysis

On the basis of this literature review, we can draw the following conclusions:

a) learning, and especially learning philosophy (or any other 'abstract' domain), is a subtle process which is inextricably linked to the autonomous (i.e. self-regulated, in a *constructivist* sense) creation and usage of *structures* (narrative, dialogical, argumentative).

b) There is an emergent trend towards the *usage* (intended in a broad sense, i.e. *production, retrieval, reuse* etc.) of structured data on the web, which can be well exemplified by the large variety of activities clustered under the 'Semantic Web' research area.

c) Despite the novelty of the area, there are already many attempts to use Semantic Web technologies for supporting learning. As we have seen, the research spectrum hosts quite different approaches; still it is possible to characterize all of them by saying that they want to create, re-use and in general benefit from the growing number of structured information repositories.

d) Because of the constantly growing availability of structured data-sources, there is a quite important SW research strand that is focusing on advanced techniques for *navigation* and *presentation* of semantic data. In particular, we have shown how such *structured* navigation mechanisms could provide features that make them apt for being used also in a learning scenario.

These are important conclusions we are able to draw from the analysis of recent research. However, now we should consider once again the motivation of our work, as presented in chapter 1. We started out by asking how the explosion of web-resources could benefit a philosophy scholar, and in particular a philosophy learner. The literature review has given us many hints suggesting how this could happen, thanks to a variety of technologies and approaches, but still a number of questions remained unanswered.

In other words, by reviewing the existing research we can see that, with respect to our initial aim, some **gaps** do exist in the literature. More precisely:

1) **Narratives for philosophy**. As we have seen, the 'digital narratives' approach (section 2.2.5) allows the dynamic re-construction of story-items into complex narrative structures. Most of the times, this process is theoretically grounded on structuralist models of *novels* (section 2.2.4). However, what are the key characteristics of the *stories* we normally *tell* in philosophy? How can the digital narrative approach be instantiated in the philosophical domain?
2) **Semantic navigations for learners**. As attested, there are various mechanisms which can be re-conducted to the concept of 'semantic navigation'; nonetheless, we believe little research investigates the employment of these technologies in an educational scenario. In other words, is a *semantic* browsing all that a learner needs, in order to start making sense of a subject domain? Are more complex or domain-specific browsing structures needed?

3) **Semantic Web in humanities' domains.** Clearly, from our investigation it appears that Semantic Web technologies (and particularly, SW technologies in e-learning) are being tested in both scientific domains, such as computer science, physics, biology, mathematics, and humanities' ones such as history of art, classic literature, history, etc. However, at a deeper level of analysis we observed that since scientific domains are traditionally highly structured they can more easily be mapped into ontologies (or, in general, formal conceptual schemas) so to be used in SW applications - e.g., think about a *gene* ontology, or an ontology of hardware components. This is not the case, instead, for humanities domains: traditionally, scholars in these areas value processes like the *subjective interpretation* and *debate* on humanities' resources, rather than the search for *objective schemas* or *taxonomies*. In other words, we realized that the task of modeling humanities domain through formal languages (so to allow computability) presents various problems which are still to be tackled by existing research. Systems such as /facet (Hildebrand et al., 2006) or CultureSampo (Eero Hyvönen et al., 2007), although dealing with humanities' contents, do not explore the type of semantic 'services' humanities' scholars often employ in their daily routines. In fact, very often such systems are just using very 'shallow' semantic models (e.g., a 'person' who *created* a 'work' which *belongs-to* a 'style'), thus oversimplifying the actual discourse that makes the discipline unique.

4) **Semantic navigations for philosophy.** Finally, we can put together the points above with a conclusive question: in the Semantic Web, what

structures do we need for navigating philosophical resources, with the aim of helping a student understand the *world* of philosophy? How is it possible to unite an active learning style, the browsing of semantic repositories and the philosophical domain? For example, are technologies such as faceted browsing (section 2.4.2) supportive enough, or do we need more domain-specific navigation mechanisms?

As mentioned in point 3, in order to provide domain-specific semantic navigations in the field of philosophy we are quite likely to need a formal specification of the philosophical domain. In SW terms, we need an *ontology for philosophy*. Thus, before moving to the description of our approach (chapter 4), in the following chapter we are going to complement this literature review by going through the most important existing models which could, at least in principle, be used as *formal representations of the philosophical domain.*