Emergent Knowledge Artifacts for Supporting Trialogical E-Learning

Yannis Tzitzikas, University of Crete and Institute of Computer Science, FORTH-ICS, Greece

Vassilis Christophides,University of Crete and Institute of Computer Science, FORTH-ICS, Greece

Giorgos Flouris, Institute of Computer Science, FORTH-ICS, Greece

Dimitris Kotzinos, University of Crete and Institute of Computer Science, FORTH-ICS, Greece

Hannu Markkanen, EVTEK University of Applied Sciences, Finland

Dimitris Plexousakis, University of Crete and Institute of Computer Science, FORTH-ICS, Greece

Nicolas Spyratos, Universite de Paris-Sud, France

ABSTRACT

This article elaborates on scenarios for collaborative knowledge creation in the spirit of the trialogical learning paradigm. According to these scenarios, the group knowledge base is formed by combining the knowledge bases of the participants, according to various methods. The provision of flexible methods for defining various aspects of the group knowledge is expected to enhance synergy in the knowledge creation process and could lead to the development of tools that overcome the inelasticities of the current knowledge creation practices. Subsequently, these scenarios are projected to various knowledge representation frameworks and for each one of them, we analyze and discuss related techniques and identify issues that are worth further research.

Keywords: *collaborative learning; electronic learning (e-learning), educational; technology; technology mediated learning; knowledge artifacts; knowledge models; semantic models; knowledge management*

INTRODUCTION

Classical learning theories are based either on the *knowledge acquisition* metaphor (where a learner individually internalizes a body of knowledge) or on the *social participation* metaphor (where a group of learners collaboratively appropriate a body of knowledge). Although widely accepted, these theories do not sufficiently capture innovative practices of both learning and working with knowledge (i.e., knowledge practices). Only sharing of knowledge in action, that is, sharing the process of learning itself, is a reliable base for developing a shared cognition (seen both as a group and an individual characteristic). In this context, the emerging theory of "*trialogical learning*" (*TL*) focuses on the social processes by which learners collectively enrich/transform their individual and shared cognition.

According to trialogical learning, knowledge creation activities rely heavily on the use, manipulation, and evolution of shared knowledge artifacts externalizing a body of (tacit or explicit) knowledge (Paavola, Lipponen, & Hakkarainen, 2004). By representing their cognitive structures or knowledge practices under the form of artifacts, individual learners can interact with themselves as well as with external tools (e.g., computers, information resources) to negotiate the meaning of concepts and signs embodied in these artifacts; this would ultimately allow them to reach a common understanding of the problem at hand. We could therefore consider the notion of shared objects of activity as the cornerstone of trialogical learning, a notion that is general enough to accommodate the requirements of various application contexts.

Shared knowledge artifacts are very useful in many applications involving some kind of collaboration. For instance, a video that records how group members carry out their tasks could be considered as a shared knowledge artifact that the group could annotate (with free text or with respect to an ontology), analyze, and further discuss (e.g., for capturing tacit group knowledge). Moreover, and more interestingly, a knowledge artifact could take a more formal substance (e.g., for capturing explicit group knowledge) as in the case of documents (e.g., a survey paper), conceptualizations (e.g., a data/knowledge base), or even software code exchanged within a group. Hereafter, we shall use the term *knowledge artifact* to refer to what is being created and/or shared by a group of learners (which could be a set of words, documents, concept maps, ontologies, annotations, etc.). It is worth mentioning that the paradigm of trialogical e-learning can be very useful within communities of practice (CoPs), as it can facilitate the negotiation of meaning and it can contribute to the development of explicit and innovative knowledge inside a CoP (Domingue, Motta, Shum, Vargas-Vera, Kalfoglou, & Farnes, 2001).

In order to communicate and meaningfully interpret their individual viewpoints, cooperating learners need to agree on a common conceptual frame of reference. Models and techniques that allow diversification and flexible amalgamation of different worldviews are still in their infancy. In this paper, we investigate various ways to build emerging knowledge spaces using the trialogical learning paradigm for eliciting the functional requirements. In particular, we focus on the various methods that can be used in order to form the common knowledge of a group by combining the individual knowledge of its members. The provision of flexible methods for defining various aspects of the group knowledge is expected to foster knowledge creation processes, and could lead to the development of tools that overcome the inelasticities of the current knowledge creation practices.

The rest of this paper is organized as follows: initially, a motivating trialogical learning scenario for collaborative knowledge creation is described and the underlying principles and interactions are presented. Subsequently, a number of methods for building emerging knowledge artifacts from individual group knowledge (of various forms) are presented, and the related knowledge management requirements are identified. Finally, the paper comments on related work and concludes.

MOTIVATING SCENARIO FOR TRIALOGICAL LEARNING

Collaborative Literature Review and Annotation

The following general scenario will be useful for the rest of this paper. Suppose that a set of N research papers, say $P=\{p_1,\ldots,p_N\}$, is given to a set of K learners $A=\{a_1,\ldots,a_K\}$ who could be students, researchers, or coworkers in a company. The goal of this group is to understand the topics discussed in these papers and to build an ontology, say O, that represents the main issues discussed in these papers. Moreover, the group has to annotate these N papers according to the derived ontology, that is, specify d(p) for each $p \in P$ where d(p) denotes the description of p with respect to O. We could also assume that there is an additional constraint, saying that the ontology should not have more than C concepts. The learners, hereafter *actors*, have to collaborate (synchronously or asynchronously) in order to carry out this task.

Note that various combinations of (N,K,C) values describe different real-life scenarios. For instance, (50,1,20) could describe what an MSc student should do in order to write the state-of-the-art of his/her MSc thesis. Of course, this scenario does not fall into trialogical learning, but is rather an instance of monological learning (knowledge acquisition metaphor). Values like (150,2,50) might describe the collaboration between a professor and a graduate student for finding a topic for a PhD thesis. Values like (100,10,10) may describe a group (comprising of 10 members) of a research lab that is trying to join a research area by studying the 100 related papers that have been published the last 5 years in the field and trying to identify the 10 main topics of the area (subsequently each member of the group would be responsible for one topic). Finally, big values for K, say 1000, could model the effort for developing an international standard.

Evaluating the Quality of Collaboratively Developed Ontologies

A related rising question is whether the "quality" of the result of this collaboration (i.e., of O and d(p)'s) could be measured and if yes, how. We can identify two broad cases, namely having an external (human or machine) observer who can grade the result or having no external grading party.

For the first case, we may, for instance, assume that there is a certain "solution" ontology (ideal or criterion), denoted by $O^{(i)}$ that is not known by the members of the group. The ontology $O^{(i)}$ could have been provided by a tutor, if there is one, or acquired via some other knowledgeable source (human or machine). Alternatively, the tutor/source might have provided a set of admissible ontologies instead of a single one. Subsequently, appropriate metrics could be employed in order to measure, at any point in time (say, at state s_i), the "distance" between $O^{(i)}$ and O_{s_i}, so that the members of the group can have a quantitative measure to rate their progress. Of course, not only the group work but also the individual work could be checked against the "solution" ontology; recall that according to Slavin (1989) and Gokhale (1995), in order to achieve effective learning in collaborative environments, both "group goals" and "individual accountability" must exist[1].

For the second case, that is, the case where no external party is available to help in the assessment of the quality of the result of the collaboration, we could probably only measure the degree of agreement between the members of the group. This measurement could be based on the following heuristic: if O_A expresses the knowledge that all members of A accept to be correct, then the bigger O_A is, the better the result of the group collaboration (assuming there is no other constraint like C in the previous scenario), since this means that a broader agreement has been reached. In both these cases the suggested methods can capture the progress of an individual or a group and can be used as the building block for tools that can record this progress.

EMERGENT KNOWLEDGE ARTIFACTS SPACES

In this section, we discuss issues that are important for supporting the previous scenario. Initially, personal and shared knowledge artifacts are introduced and their relation is clarified. Subsequently, we show how a set of learners can interact on the basis of their personal and shared knowledge artifacts.

Personal vs. Shared Knowledge Artifacts

To abstract from representation details, we shall hereafter use the term knowledge base (KB) to refer to an ontology or to an ontology-based information base (i.e., to a set of objects annotated with ontological descriptions).

Although trialogical learning focuses on shared artifacts, learners should be able to construct and evolve their own models. Let KB_a denote the KB of an actor a. Now let KB_A denote the "shared" (or common) KB of a set of actors A. The important issue here is the relation between KB_A and KB_a (for $a \in A$). Here we identify three broad cases:

- *UNION-case*. Here KB_A is obtained by taking the union of the KBs of all participants, that is: $KB_A = \cup\{KB_a \mid a \in A\}$. Notice that different models may encode different viewpoints on a domain, so their union is not necessarily a coherent whole. Therefore, depending on the context, KB_A could be inconsistent (if there is a notion of consistency). For example, if the task is to annotate a video with argumentative maps, then consistency is not an issue. If, on the other hand, the task is to develop an ontology (for subsequently building a bibliographic database) or a software module, then consistency is a very important issue.
- *INTERSECTION-case*. Here KB_A is obtained by taking the intersection of the KBs of all participants, that is: $KB_A = \cap\{KB_a \mid a \in A\}$, so it contains statements "accepted" by all participants.
- *QUANTITATIVE-case*. Here KB_A is defined by a quantitative method, for example, it contains all sentences that are accepted by at least a percentage of the actors. Obviously, UNION and INTERSECTION are special cases of the QUANTITATIVE case.

Interaction Through Knowledge Artifacts

Suppose that we want to design and develop an application for supporting various forms of collaboration (e.g., asynchronous and synchronous) as well as personal and shared knowledge artifacts. Figure 1 sketches a possible UI[2] for such an application that could serve as a proof of concept and as a gnomon for identifying and analyzing the associated technical requirements and challenges.

The UI is divided in two main areas: the left area allows managing the personal space, while the right area allows managing the group space. The personal space (left area) is under the full control of the respective learner, so everything is editable in that area; the right area shows the shared artifacts and constitutes the key point for collaboration and for supporting trialogical e-learning. Assuming the scenario described earlier, each user may develop his or her own ontology at the left area, while the right window shows the group ontology O; the group ontology O has been derived from the personal ontologies using any of the methods mentioned in the previous subsection.

The relationship between the personal space and the group space is very important. The button labeled by "→" allows a user to copy the desired parts from his/her ontology to the group space. The button labeled by "←" allows a user to copy the desired parts from the group ontology to her personal space. An option that keeps the button "→" permanently pressed would allow synchronous collaboration in the sense that every change at a learner's ontology would be immediately reflected (propagated) to the group ontology (e.g., blackboard-based collaboration). Symmetrically, an option that keeps the button "←" permanently pressed would immediately propagate any changes on O to the personal space[3]. Deletions are handled analogously and are discussed in a subsequent section. The issue of consistency is discussed in a subsequent section. Systems (and UIs) that allow this kind of collaboration/interaction will be called *synodic*[4].

Of course, the sketched scenario (and UI) of trialogical e-learning can be enriched with a plethora of auxiliary functionalities. Next we identify the ones we consider the most important:

- The *group space* view could be *customizable*, for example, instead of showing the group ontology, one participant may want to see the ontology derived by considering the ontologies of only a *subset* of the participants. Moreover, in some cases, some additional flexibility on how the shared KB is defined could be useful; to allow this, the shared KB should be definable using any set theoretic expression over subsets of A (instead of a simple union or intersection over all elements of A). For example, $K_{(\{a_1\}\cap\{a_2\})\cup(\{a_3\}\cap\{a_4\})}$ could capture the scenario where two groups a_1, a_2 and a_3, a_4 collaborate in the sense that the joint work of each group is integrated. Optionally, the group space could be *managed* by a person whose role would be to accept or reject the changes that the participants forward to the group ontology.
- In many cases, it is important to be able to access the *provenance* of a statement. For this reason, provenance information should be recorded and be available at any time to interested participants. Moreover, the participants should be able to *annotate* any element of their personal or group space. Such

Figure 1: An indicative UI for trialogical e-learning

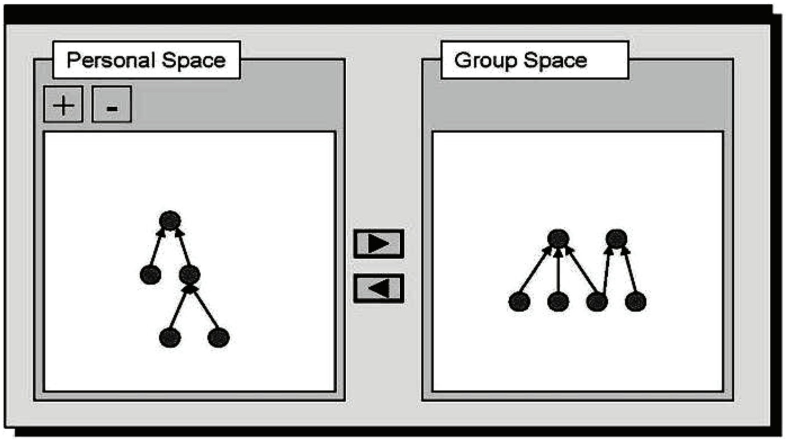

annotations could be textual or ontology-based.

- *Usability* is always a very important issue. For instance, by placing the mouse on top of an element of the group ontology, a balloon should open showing who provided this info (and/or what percent of the actors agree with this). Moreover, the *visualization* of knowledge artifacts is a very important, challenging, and open issue; a brief related discussion can be found in Liebig and Noppens (2005) and Tzitzikas and Hainaut (2006).
- The UI could be enriched with *teleconferencing* services allowing the participants to discuss in real-time while using the system.

SYNTHESIZING KBS

To support trialogical e-learning in our motivating scenario, we need to support the formation and evolution of A, P, O, and d(p)'s. In order to identify the distinctive knowledge management requirements for this task, we will first present a refined approach for supporting personal and shared knowledge artifacts, and then we will investigate the effects of using various forms of KBs, starting from very simple ones and gradually considering more complex KB forms. The reason for trying to identify the key knowledge management requirements (that originate from trialogical learning) is to investigate how we could support them by extending, accordingly, the core knowledge management technologies (instead of developing yet another e-learning application).

Supporting Personal and Shared Knowledge Artifacts

The personal space of an actor could be divided into two spaces: one *private* and one *public*. The group (shared) space is derived from the public personal spaces of the actors. Similarly, each actor a_i has two unique *identifiers*: one *private* and one *public*. The first, denoted by a_i^p, is associated with every "statement" (construct or update operation) concerning his personal space (i.e., the statements in the private personal KB of the learner). The second, denoted by a_i, is associated to every statement he has forwarded to the group space (i.e., the statements in the public personal KB of the learner). Let KB_i^p denote the KB containing all statements with identifier a_i^p, and KB_i denote the KB containing all statements with identifier a_i. Normally, it should be the case that $KB_i \subseteq KB_i^p$, that is, the public personal base of a user should be a subset of his personal private base. However, in social life, sometimes people forejudge or "pretend" that they accept facts although they do not really believe them (e.g., because all other people do, or for strategic reasons). In such cases the relationship $KB_i \subseteq KB_i^p$ does not hold. We understand that, in the context of communities of practice (where all members are somehow "committed" to augment the individual and organizational knowledge bases), there is a high level of trust and a strong sense of responsibility that would usually make this a nonissue. Nevertheless, we choose to be more liberal in order to cover cases where this is not true so, in order to leave learners free, we suggest that no constraint should be imposed on KB_i and KB_i^p. The important point here is that the synthesis (or amalgamation) of all KB_is forms the shared artifacts of the group (i.e., the shared artifacts according to trialogical learning). Figure 2 illustrates the idea.

KB = A Set of Words

Let us now consider that the learners' (and the shared) knowledge (KB) is just a set of words (i.e., a set of strings). In our application scenario, this corresponds to the case where the ontology that the learners have to create has the form of a set of keywords.

For reasons explained in the previous subsection, we need, for each actor $a_i \in A$, two KBs: KB_i^p and KB_i. The first (KB_i^p) is a set of pairs of the form (w,a_i^p), whereas the second (KB_i) is a set of pairs of the form (w,a_i) where w is a word. At the beginning of a learning session it could be $KB_i^p = KB_i = \emptyset$ for each i=1,…,K, although this is not a necessary constraint.

Consider now an actor a_i who uses the left area of the UI and creates a KB_i^p. Now suppose that he selects some elements of KB_i^p, say a word w, and presses the "→" button. One reaction to this event can be

1. A new pair (w,a_i) is created.
2. The group KB is updated according to this information (depending on the way that the group KB is defined).

Now suppose the user selects some elements, say a word w, from the group space (rightmost area) and presses the "←" button. One reaction to this event can be

3. A new pair (w,a_i^p) is created. This step makes the assumption that the user agrees with w. In other words, we treat this case as if the user had added the word w to his private base himself.
4. The private base of the user is updated accordingly.
5. Probably (or optionally) a pair (w,a_i) should be created.

Let us now suppose that the user *deletes* one element w from his private KB. If the user had "published" w in the past, that is, if a pair (w,a_i) exists, then the system should ask the user if the pair (w,a_i) should be also deleted from his public KB (KB_i) or not. This case suggests that it would be more informative if the UI for each actor a_i

Figure 2: The overall picture

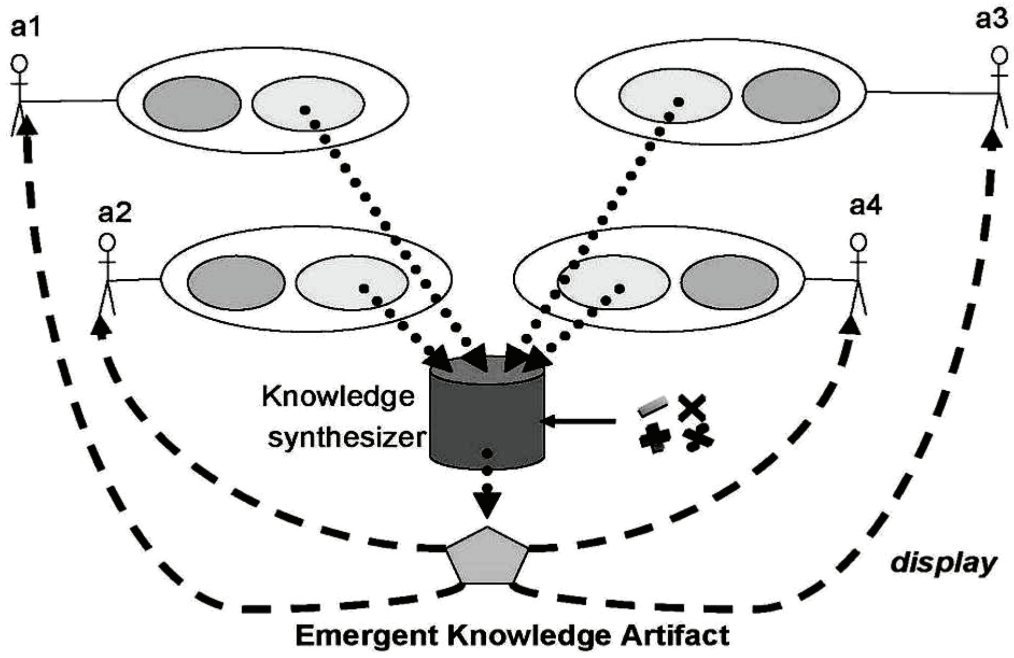

were divided into three areas, one for each of KB_i^p, KB_i, and KB_A (see Figure 3). This would allow monitoring and controlling the contents of KB_i as well. Notice that, if the UI of Figure 3 is used, all elements placed in the public KB are automatically propagated to the group KB, so a "→" symbol from the public KB to the group KB is not necessary.

Let us now investigate how the "shared" KB could be defined. Let KB_A denote the KB obtained by taking the union of the public bases of all actors, that is, $KB_A = \cup_1^K KB_i$. We can define the *support* of a word w, denoted by for(w), as the set of ids that correspond to actors who have included w in their public KB. So KB_A can also be considered as a set of pairs of the form (w,for(w)) where $for(w)=\{a_i \mid (w,a_i) \in KB_i\}$. Notice that this view is quite generic as it allows defining the group KB at run-time using various methods (union, intersection, or other); in the following cases, we denote by |S| the number of elements in a set S (so, for example, |for(w)| would denote the number of elements in the support of w, namely for(w)):

- The UNION case would include all words w such that $|for(w)| \geq 1$, specifically:
  $KB_{\cup A} = \{w \mid for(w) \subseteq A \text{ and } for(w) \neq \emptyset\}$
- The INTERSECTION-case would include all words w such that $|for(w)|=K$, specifically:
  $KB_{\cap A} = \{w \mid for(w) \supseteq A\}$
- The z-PERCENT case would include all words w such that $|for(w)|/K \geq z$, specifically:

Figure 3. A (refined) indicative UI for trialogical e-learning

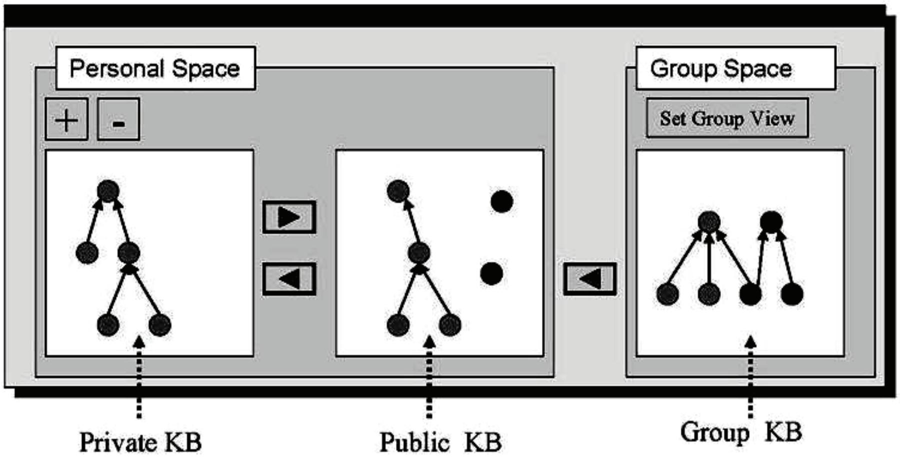

$$KB_{z\%A} = \{w \mid \frac{|for(w) \cap A|}{|A|} \geq z\}$$

- The case where a user wants to see the group ontology, as derived by considering only a subset A′ of A, can also be captured by these formulas by replacing A with A′.

It has been made evident that by considering a KB as a set of pairs of the form (w,for(w)), we can compute "whatever shared KB" we want. So such a representation could be adopted for the physical layer of the repository.

As already mentioned, a related problem is the evaluation of the learners' understanding progress of the domain concepts. Here, we will consider the case where some externally provided "correct" ontology is available, so the evaluation problem is reduced to the problem of assessing the similarity between this "target" ontology and the learners' ontologies (group, public or private). Let W and W′ be the set of words stored in KB and KB′ respectively; then, we can define the distance between KB and KB′ on the basis of W and W′ using any distance metric for sets. For instance, we can use the symmetric difference, that is, dist(KB,KB′)=|W\W′|+|W′\W|, the Dice coefficient, that is, $dist(KB, KB') = 1 - \frac{|W \cap W'|}{|W \cup W'|}$, or any other metric that we find suitable for the application at hand.

KB = A Binary Relation

Now suppose that a KB is a binary relation R over a set of elements T, ($R \subseteq T^2$). Let r denote an element of a R, for example, r=(t,t′) where t, t′∈T. In our application scenario, this corresponds to the case where the ontology (that the learners have to create) is a graph of keywords.

The personal and group KBs in this case can be defined in the same manner as in the previous subsection (e.g., for the union case: $KB_{\cup A}=\{w \mid for(w)\subseteq A$ and $for(w)\neq\emptyset\}$). The only difference is whether the set T is considered to be known by all actors (and thus is not part of the created knowledge), or not. If T is considered part of the created knowledge, then the KB of an actor could be characterized by both R_i and T_i (of course $R_i\subseteq T_i^2$); therefore, we can define shared KBs (e.g., $KB_{\cup A}$ or $KB_{\cap A}$) not only for R but also for T.

KB = A Binary Relation with Second Order Properties

Here we consider again the case where a KB is a binary relation R over a set of elements T ($R\subseteq T^2$), with the extra constraint that this relation satisfies some properties (for example, R could be constrained to be reflexive, symmetric, and transitive). These extra properties can be seen as derivation rules (inferences) or constraints and allow us to capture more interesting cases. For instance, assuming that R is a preorder (i.e., a reflexive and transitive relation) allows us to capture the case of taxonomies; in our application scenario, this corresponds to the case where the ontology (that the learners have to create) has the form of a taxonomy. Therefore, supporting this scenario allows us (among other things) to support collaborative (and trialogical) taxonomy construction.

We could model these derivation rules (e.g., transitivity) by defining a *consequence operator* Cons that models inference services. Considering a KB as a set of sentences S, the consequence operator returns all the sentences that can be inferred by S (obviously, it should hold that $S\subseteq Cons(S)$). Alternatively, axioms could be modeled using the notion of consistency.

The introduction of Cons allows us to consider, apart from KB_i and KB_i^p, the sets $Cons(KB_i)$ and $Cons(KB_i^p)$ (respectively) as well, for each i=1,…,K. Notice that the consequences of a KB form a different (richer) set than the KB itself; thus, defining a shared KB on the basis of $Cons(KB_i)$ would (in general) give different results than defining it on the basis of KB_i and either of the KBs could be used for the definition of the "shared" KB. This fact is illustrated in the example of Figure 4, where $KB_{\cap\{1*2*\}}$ has been used to denote that $Cons(KB_1)$ and $Cons(KB_2)$ were used for the definition of $KB_{\cap\{1,2\}}$.

KB = A Total Order

This case is actually a special case of this subsection, so the same general comments apply. However, the case where R is a total order is particularly interesting, so we chose to study it separately. Total orders are useful, for instance, when learners have to rank a set of available options T in order to come up with some decision, such as ranking a set of keywords or a set of papers according to their significance or importance; the latter case appears in the selection process of peer-reviewed scientific conferences and journals, where the shared (group) KB can be obtained by aggregating the "rankings" of the learners (reviewers). Total orders are also useful for modeling the case where a questionnaire comprising multiple choice questions (where more than one choices are correct for each question) is presented to the learners and the tutor asks from the group to mark only one choice (the most appropriate).

The problem of aggregating the personal total orders to generate a commonly agreed upon total order has been extensively studied in the literature and could be directly used for our purposes, that is, the aggregation of personal KBs (total orders) to gener-

Figure 4. Local reasoning and group KBs

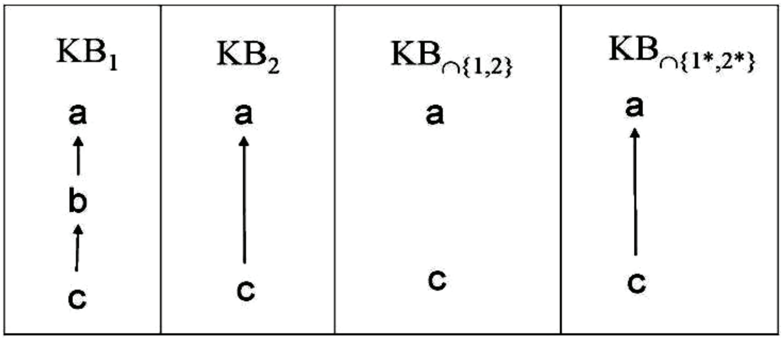

ate a group KB (total order). For example, we could adopt various techniques (mainly coming from the area of Social Choice), like plurality ranking, Borda ranking (Borda, 1781), Condorcet ranking (Condorcet, 1785), or Kemeny Optimal Aggregation (Kemeny, 1959), but we should not forget the Arrow's impossibility theorem (Arrow, 1951). A Borda-like technique for aggregating weakly ordered subsets of a set that could be used for our purposes is described in (Tzitzikas, 2001)

Collaborative selection and filtering (i.e., the provision of prediction and recommendation services) is also related to this case (and also useful for collaborative knowledge creation and learning). The difference with the standard total order case is that now actors do not rank a set of objects, but they rate (using a numerical scale) a subset of the objects (e.g., instead of rankings of the form $<o_1,o_2,o_3>$ meaning that o_1 is preferable to o_2 which is preferable to o_3, we may have input of the form {score(o_1)=5, score(o_2)=3, score(o_3)=2}). Furthermore, it is worth to investigate generalizing these techniques for the case where, instead of atomic objects, we have structured knowledge artifacts (e.g., a conceptual graph expressed in RDF).

In this scenario, the set T is not part of the created knowledge (in other words, it preexists); this need not be the case. Suppose that a group of persons (e.g., the authors of the current paper) would like to collaborate in order to specify the structure of a research paper to be submitted to IJWLTT. Each one proposes a structure, that is, a total order of strings (here a string can be the title of a section or a short paragraph indicating the contents that this section should have). The collaborative system should aid them to come up with some decision, that is, with one structure either accepted by all of them or by most of them. As it would not be realistic to expect that two persons will propose exactly the same title (or paragraph) for a section, a text similarity function could be employed (meaning that two texts with degree of similarity greater than a certain threshold could be considered to denote the same section). As each participant will be able to see what the others do (using the right area of the UI), they are expected to refine, improve, or change the pieces of text they have provided (and their relative order) while interacting with the system. After some interactions the group will hopefully reach a structure that is probably better than what each one could do by him- or herself (of course apostates may exist).

An alternative method to support this scenario follows. Suppose that the paper to be submitted should have exactly seven sections. Let T be the pieces of texts that all actors have provided (i.e., $T=\cup_1^K T_i$), for example, if K=3 then $|T|\leq 21$. The group KB (group paper structure) could be the result of applying the K-Means clustering algorithm (here K=7, so we have 7-Means) on T, resulting to a set T_A (each element of T_A would be a set of texts). The ordering of the elements of T_A could be derived by first mapping the participant's rankings to rankings of T_A and then applying a rank aggregation method. We have just described a collaborative (or cooperative) document-authoring scenario.

KB = An RDF-Based Repository

Suppose now the case where the learners have to create an ontology-based repository (ontology plus ontology-based metadata) using the RDF language. A repository of this kind has the form of a conceptual graph. According to RDF (Brickley & Guha, 1999; Miller & Swick, 2003), this graph can be seen as a set of RDF triples that actually defines a directed graph consisting of three kinds of relations (instanceOf, isA and property). Therefore, we could write $KB=(R_{in},R_{isa},R_p)$, where R_{in} contains all instanceOf relationships, R_{isa} contains all isA relationships, and R_p contains all property relationships. Note that the *isA* relation (R_{isa}) models a transitive relation, so the issues discussed in the subsection about second-order properties apply here as well. It follows that the semantics of the RDF constructs should be taken into account when applying operations (i.e., union and intersection) on various KBs; such issues for RDF are discussed in Kaoudi, Dalamagas, and Sellis (2005).

The increased complexity of RDF (with respect to the other cases handled in this paper) gives rise to certain issues, such as the issue of consistency and inconsistency. If inconsistency arises in one individual (personal) KB, then the user is responsible for making what is necessary for reaching a consistent one[5]. However, one can easily see that even if each individual personal KB is consistent, this is by no means a guarantee that the group ontology will be also consistent. So, what should we do when faced with an inconsistent group KB (ontology)? Who and how should react in that case? Should the system allow such cases? Is there anything it could do for aiding actors to overcome this problem?

Our opinion is that it would not be flexible to forbid inconsistent group KBs; the system should allow inconsistent group KBs but, at the same time, it should be able to detect and highlight inconsistencies, allowing the actors (learners) to deal with the problem (if necessary). Allowing inconsistency in the personal KBs as well, gives rise to another interesting case: the individual KBs could be inconsistent while the group KB is consistent[6].

For tackling inconsistency at the group level, a powerful knowledge manager could try to derive (and present) consistent subsets of the group KB. It could also probably adopt a quantitative notion of consistency (instead of the dichotomy of KBs to consistent and inconsistent). Let us use the notation $\models$KB to denote that KB is consistent. If a KB is inconsistent ($\not\models$KB), then the system could try computing $KB_{A'}$ (specifically, $KB_{\cup A'}$, or $KB_{\cap A'}$, or $KB_{z\%A'}$), where A' is the maximal subset A' of A such that $\models KB_{A'}$ (resp. $\models KB_{\cup A'}$, or $\models KB_{\cap A'}$, or $\models KB_{z\%A'}$). Notice that if there is no inconsistency, the definitions of group KBs coincide with the original ones.

A more sophisticated method would be to define a notion of *ranking* (or *priority*) that could be attached to each RDF triple in the repository. This ranking would encode the relative strength (reliability) of each triple as per each individual learner's understanding and could be either qualitative (i.e., encode the ranking through a full or partial order) or quantitative (i.e., encode the ranking through a numerical assignment of a priority to each triple, which implies an ordering). This refinement facilitates the definition of a quantitative notion of inconsistency, as well as the process of aggregation using techniques from Social Choice, described previously. Furthermore, it allows the adaptation of works related to belief merging (Konieczny, 2004; Konieczny, Lang & Marquis, 2004; Konieczny & Perez 2005) in our aggregation context, by facilitating the formal description of notions like "weakening," "conceding," and "negotiating" (Konieczny, 2004), the development of *arbitration* or *majority* merging operators (Konieczny & Perez, 2005), and the definition of distances and aggregation functions (Konieczny et al., 2004).

Notice that, unlike traditional approaches that conceive ontologies as thorough engineering artifacts issued by strict design processes and policies, in trialogical learning, the ontology creation and evolution can be seen as a social process where learners collectively improve their individual and shared understanding through social interaction, something that pertains to all described cases so far. In this context, the individual interactions of group members would lead to global effects that could be observed as *emerging* knowledge artifacts (related somehow to emergent semantics (Aberer, Catarci, Cudre-Mauroux, Dillon, Grimm, Hacid, Illarramendi, et al., 2004)). Ontologies would thus become an emergent effect of open-ended interactions within or across groups of individuals, as opposed to a firm commitment of a small group of domain experts (for more see Mika, 2005).

Finally, we should remark that workflow issues are orthogonal to the issues we discussed so far, because the issues elaborated so far arise in almost every step of any workflow process that should be carried out collaboratively.

RELATED WORK

In this section, we discuss a number of further issues and related work. There is a plethora of works (both theoretical and applied[7]) related to collaborative knowledge construction. Although the majority of them were not proposed (or applied) for e-learning, they are related to the main theme of this paper (trialogical learning).

Advancing the Expressive Power of the Representation Frameworks

One particularly interesting related class of works deals with issues that arise as we step up the expressive power of the representation framework:

- Knowledge change and evolution raises various issues like the distinction between *update* and *revision* (in the sense defined in Katsuno & Mendelzon, 1991) or the applicability of belief revision theories to ontology evolution (Flouris, 2006).

- Measuring the distance between two KBs (e.g., for grading as described in a previous section) may not be enough by itself; for learning purposes, it might be also important to compute and show the difference, or *delta*, between two KBs. Some approaches for computing deltas of RDF graphs are described in SemVersion (Volkel, Winkler, Sure, Kruk, & Synak, 2005), PromptDiff (Noy, Kunnatur, Klein, & Musen, 2004), and (Berners-Lee & Connolly, 2004).

Increasing the Number of Actors

As the number of actors scales up, additional issues arise, such as the need for social network analysis. It is worth mentioning here that the Web is probably a case of collaborative knowledge creation of a very primitive form. The actors of the Web can only create and update their own KB_i's (interlinked Web pages) and the only method to combine the KBs of different actors is to add one-way links between them. Despite this simplicity, the growth of the Web was (and remains to be) astonishing, mainly because no one ever tried to impose a structure or any form of control on that. Therefore, link analysis techniques (either applied on social networks, or on articulated KBs (Brin & Page, 1998; Guo, Shao, Botev, & Shanmugasundaram, 2003), or on large KBs (Tzitzikas & Hainaut, 2005)) are also expected to be useful in large-scale collaborative knowledge creation.

An architecture for building KBs by mass collaboration where machine-learning techniques are employed for estimating the quality of knowledge is described in Richardson and Domingos (2003). The quality of the knowledge contributed by volunteers is also the main theme of Chklovski and Gil (2005), whereas some general ideas about collaborative knowledge construction (e.g., assigning values and credits to knowledge contributors) are discussed in Martin, Blumenstein, and Deer (2005). These works mainly concern the case where the number of participants is very high.

The provision of personalized services is very useful in large-sized KBs as well (Spyratos & Christophides, 2005). Mass collaboration and provision of personalized services is the main motivation behind folksonomies (Mathes, 2004; Mika, 2005; Ohmukai, Hamasaki, & Takeda, 2005), so our proposal could be placed in the context of the work described in Mika (2005).

Modular Knowledge Spaces

The need for defining separate knowledge spaces has been identified in several contexts (including the semantic Web) as this would be useful for data syndication, for restricting information usage, and for access control, among others. Several approaches have been proposed (like Euzenat, 1996; Guha, McCool, & Fikes, 2004; Jannink, Pichai, Verheijen, & Wiederhold, 1998;Wiederhold, 1994), the more recent one being that of *named graphs* (Carroll, Bizer, Hayes, & Stickler, 2005; Watkins & Nicole, 2006). Along the same direction, in Bao, Caragea, & Honavar (2006), packaged-extended description logics, as well as semantics for collaborative ontology construction, are proposed. Our paper goes one step further by stressing the need for synthesizing such knowledge spaces.

Combining KBs

The need for combining KBs has also been identified in several contexts. There are some theoretical works (e.g., Konieczny, 2000; Konieczny, 2004; Konieczny, Lang, & Marquis, 2004; Konieczny & Perez, 2005), coming from the area of belief merging, that describe operators for combining KBs. For instance, Konieczny (2000) studies the properties of a number of merging operators where a KB is a set of first-order formulae. Although very interesting theoretically, the practical exploitation of these results is quite distant.

Semantic Wikis

Wikis can serve as the knowledge platform for a community of practice, where its members can interact with each other, share information and knowledge, and discuss about them. Semantic Wikis (Aumueller & Auer, 2005; Krotzsch, Vrandecic, & Volkel, 2005; Tazzoli, Castagna, & Campanini, 2004; Volkel & Oren, 2005) can be used to support knowledge management and learning activities of communities of practice; these activities can be seen as a collaborative learning effort being supported by the collaborative features (editing, versioning, discussion) of the semantic wikis. Actually, wikis are already used in teaching (Lamb, 2004) under different contexts and for different purposes. Wikis can support both informal and formal learning by stimulating learners through their ability to invite them to respond immediately to posted problems and questions, either by themselves or through collaboration with others.

Apart from better structuring content (through semantic Web technologies) in order to make it understandable by the machines, semantic wikis can also be used as a collaborative environment for creating or evolving ontologies, like in Hoehndorf, Prufer, Backhaus, Visagie, and Kelso (2006). In these efforts, the main interest is in providing the means for collaborative editing of ontologies through the provided wiki functionalities. Usually this functionality works well for "light-weight" ontologies; however, when it comes to "heavy-weight" ontologies certain problems arise, because the process of collaborative editing has not been formally described (Corcho, Fernandez-Lopez, & Gomez-Perez, 2003). Nevertheless, semantic wikis support the collaborative knowledge formation process (Schaffert, Gruber, & Westenthaler, 2005), which is valuable in communities of practice and knowledge engineering in general.

Ontologies and Collaborative Knowledge Creation

This section describes, in brief, three indicative related systems.

The Ontolingua Server (Farquhar, Fikes, & Rice, 1997) uses a notion of users and groups that is typical in most multiuser file systems. The server provides support for simultaneous work through group sessions using access control mechanisms. There is a notion of group ownership, which enables other members of that group to join the session and work simultaneously on the same set of ontologies. In addition, a notification mechanism is supported that informs each user of other users' changes (in terms of basic operations such as add, delete, and modify).

The Adaptive Presentation Environment for Collaborative Knowledge Structuring (APECKS) system (Tennison & Shadbolt, 1998) is an ontology server supporting collaboration by allowing domain experts to create ontologies based on their own perspective. APECKS allows perspectives' comparison and discussion on the differences found: the emphasis in APECKS is not on the outcome, but rather on the process of development itself, that is, the disagreements and discussion required in order to create a consensual ontology.

The Collaborative Construction of Consensual Knowledge (CO4) system

(Euzenat, 1995) is designed for the incremental and concurrent building of a KB. The implementation of the CO4 system supports collaborative construction of a formal KB and addresses the problem of consensus of the KB with the help of the CO4 protocol for integrating knowledge. The principles underlying the CO4 protocol are derived from the peer-reviewing protocol (Peters, 1995): before being committed into a consensual KB, the knowledge must be submitted, reviewed, and accepted by the community; this way, at the end of the development process, knowledge stored in the KB is safe enough so that anybody can accept it and use it confidently and easily.

Current E-Learning Approaches and Standards

Current e-learning systems are based on the notion of learning objects (LOs). LOs capture any chunk of learning material regardless of its form, granularity, and functionality, so, by definition, they encapsulate both learning content and appropriate descriptive information (i.e., metadata). Several e-learning specifications (based on LOs) have been proposed in the literature, like ARIADNE (Foundation, 2006), IMS (IMS, 2002), and LOM (IEEE, 2002); these have been recently encoded using semantic Web languages like RDF/S (McGreal & Roberts, 2001; W3C, 2004).

Learning management systems (LMSs) and the reference models they implement, like the shareable content object reference model (SCORM) and the open knowledge initiative (OKI), provide frameworks and object models that allow for a more structured view and development. For instance, SCORM supports sequencing rules that can be used to specify learning paths of individual LOs rather than semantic relationships of LOs committing to e-learning RDF/S schemas. That way, the association among different LOs takes place only at design/authoring time, not allowing for on-demand browsing of LOs; thus, it cannot take advantage of associations established by the learners or other instructors. A work that allows on-demand sequencing of LOs and the creation of learning paths (shared across the learners) is described in Kotzinos, Pediaditaki, Apostolidis, Athanasis, and Christophides (2005). It follows that flexible collaborative knowledge creation services (in the form prescribed in this paper) are not currently supported by LMSs.

EPILOGUE

This paper described a specific scenario for collaborative knowledge creation in the spirit of the trialogical learning paradigm. According to this scenario, the group KB is formed by combining the KBs of the participants, using various methods. The provision of flexible methods for defining various aspects of the group knowledge is expected to enhance synergy in the knowledge creation process and could lead to the development of tools that overcome the inelasticities of the current knowledge creation practices. An indicative UI was sketched enabling us to scent the most important issues that are raised for its realization. Subsequently, we focused on knowledge management and projected this scenario to various knowledge representation frameworks, and for each one we outlined related application scenarios, techniques, and issues that are worth further research.

Trialogical e-learning requires advanced knowledge management services, probably more advanced than those that have emerged in the database and KR area (including the semantic Web). Database and KR technologies have provided stable solutions for the case where there is a commonly accepted conceptualization and worldview, but methodologies and technologies that allow diversification and flexible amalgamation of different worldviews (a necessary feature for trialogical learning) are still in their infancy. Areas of knowledge management that are related (in principle) to trialogical e-learning include modal logics, quantitative methods for aggregating knowledge, and belief revision theories.

We are currently investigating and experimenting with these issues in the context of the Knowledge Practices Laboratory (KP-Lab) project (cofunded by the IST programme of the EU 6). The implementation will be based on semantic Web technologies, namely on the RDF Suite (FORTH-ICS, 2005; Karvounarakis, Christophides, & Plexousakis, 2002; Magiridou, Sahtouris, Christophides, & Koubarakis, 2005).

ACKNOWLEDGMENTS

This research was conducted within the Knowledge Practices Laboratory (KP-Lab) Integrated Project sponsored under the 6th EU Framework Programme for Research and Development. The authors are solely responsible for the content of this article. It does not represent the opinion of the KP-Lab consortium or the European Community, and the European Community is not responsible for any use that might be made of data appearing therein.

REFERENCES

Aberer, K., Catarci, T., Cudre-Mauroux, P., Dillon, T. S., Grimm, S., Hacid, M. S., Illarramendi, A., Jarrar, M., Kashyap, V., Mecella, M., Mena, E., Neuhold, E.J., Ouksel, A. M., Risse, T., Scannapieco, M., Saltor, F., de Santis, L., Spaccapietra, S., Staab, S., Studer, R., & De Troyer, O. (2004). Emergent semantics systems. *In Proceedings of the 1st International IFIP Conference on Semantics of a Networked World (ICSNW-04)*, (pp. 14-43).

Arrow, K. J. (1951). *Social choice and individual values*. New York: Wiley, 2nd ed. 1963.

Aumueller, D., & Auer, S. (2005). Towards a semantic wiki experience-desktop integration and interactivity in Wiksar. In *Proceedings of Semantic Desktop Workshop*, at ISWC-05, Galway, Ireland.

Bao, J., Caragea, D., & Honavar, V. (2006). Towards collaborative environments for ontology construction and sharing. In *Proceedings of the International Symposium on Collaborative Technologies and Systems (CTS-06)*, (pp. 99-108), Washington DC, IEEE Computer Society.

Berners-Lee, T., & Connolly, D. (2004). *Delta: An ontology for the distribution of differences between RDF graphs*. Retrieved May 1, 2004, from http://www.w3.org/DesignIssues/Diff

Borda, J. de. (1781). *Memoire sur les Elections au Scrutin*. Paris: Histoire de l'Academie Royale des Sciences.

Brickley, D., & Guha, R. V. (1999). *Resource description framework (RDF) schema specification: Proposed recommendation.* Retrieved January 24, 2007, from http://www.w3.org/TR/1999/PR-rdf-schema-19990303

Brin, S., & Page, L. (1998). The anatomy of a large-scale hypertextual Web search engine. In *Proceedings of the 7th International WWW Conference (WWW-98)*. Brisbane, Australia.

Carroll, J., Bizer, C., Hayes, P., & Stickler, P. (2005). Named graphs, provenance and trust. In *Proceedings of the 14th International WWW Conference (WWW-05)*, (pp.

613-622).

Chklovski, T., & Gil, Y. (2005). Towards managing knowledge collection from volunteer contributors. In *Proceedings of the 2005 AAAI Spring Symposium on Knowledge Collection from Volunteer Contributors*.

Condorcet, M. de. (1785). *Essai sur l'Application de l'Analyse a la Probabilite des Decisions Rendues a la Pluralite des Voix*. Paris.

Corcho, O., Fernandez-Lopez, M., & Gomez-Perez, A. (2003). Methodologies, tools and languages for building ontologies: Where is their meeting point? *Data & Knowledge Engineering*, *46*(1)*,* 41-64.

Domingue, J., Motta, E., Shum, S. B., Vargas-Vera, M., Kalfoglou, Y., & Farnes, N. (2001). Supporting ontology driven document enrichment within communities of practice. In *Proceedings of the 1st International Conference on Knowledge Capture (K-CAP-01)*, (pp. 30-37). New York: ACM Press.

Euzenat, J. (1995). Building consensual knowledge bases: Context and architecture. In *Proceedings of the 2nd International Conference on Building and Sharing Very Large-Scale Knowledge Bases (KBKS-95)*.

Euzenat, J. (1996).Corporate memory through cooperative creation of knowledge based and hyper-documents. In *Proceedings of the 10th Knowledge Acquisition Workshop (KAW-96)*.

Farquhar, A., Fikes, R., & Rice, J. (1997). The ontolingua server: A tool for collaborative ontology construction. *International Journal of Human Computer Studies*, *46*(6)*,* 707-727.

Flouris, G. (2006). *On belief change and ontology evolution*. Doctoral Dissertation, Computer Science Department, University of Crete, Greece.

FORTH-ICS. (2005). The ICS-FORTH RDF-suite: High-level scalable tools for the semantic Web. Retrieved January 24, 2007, from http://139.91.183.30:9090/RDF

Foundation, A. (2006). ARIADNE. Retrieved January 24, 2007, from http://www.ariadne-eu.org/

Gokhale, A. A. (1995). Collaborative learning enhances critical thinking. *Journal of Technology Education*, *7*(1).

Guha, R., McCool, R., & Fikes, R. (2004). Contexts for the semantic Web. In *Proceedings of the 3rd International Semantic Web Conference (ISWC-04)*, Hiroshima, Japan.

Guo, L., Shao, F., Botev, C., & Shanmugasundaram, J. (2003). XRANK: Ranked keyword search over XML documents. In *International Conference on Management of Data, SIGMOD 2003*, San Diego, California.

Hoehndorf, R., Prufer, K., Backhaus, M., Visagie, J., & Kelso, J. (2006). *The design of a wiki-based curation system for the ontology of functions*. Retrieved January 24, 2007, from http://bio-ontologies.man.ac.uk/download/HoehndorfEtAlWiki.pdf

IEEE (2002). *Learning object metadata*. Retrieved January 24, 2007, from http://ltsc.ieee.org/doc/wg12/LOM_WD6_4.pdf

IMS (2002). *Content packaging information model*. Retrieved January 24, 2007, from http://www.imsproject.org/content/packaging/

Jannink, J., Pichai, S., Verheijen, D., & Wiederhold, G. (1998). Encapsulation and composition of ontologies. In *Proceedings of the 1998 AAAI Workshop on AI and Information Integration*.

Kaoudi, Z., Dalamagas, T., & Sellis, T. (2005). RDFSculpt: Managing RDF schemas under set-like semantics. In *Proceedings of the 2nd European Semantic Web Conference*

(ESWC-05), Crete, Greece.

Karacapilidis, N., & Papadias, D. (2001). Computer supported argumentation and collaborative decision making: The HERMES system. *Information Systems, 26*(4), 259-277.

Karvounarakis, G., Christophides, V., & Plexousakis, D. (2002). RQL: A declarative query language for RDF. In *Proceedings of the 11th International WWW Conference (WWW-02)*, Hawaii, USA.

Katsuno, H., & Mendelzon, A. O. (1991). On the difference between updating a knowledge base and revising it. In *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR-91)*, (pp. 380-395).

Kemeny, J. (1959). Mathematics without numbers. *Daedalus*, *88*, 571-591.

Konieczny, S. (2000). On the difference between merging knowledge bases and combining them. In *Proceedings of the 7th International Conference on Principles of Knowledge Representation and Reasoning (KR-00)*, 135-144.

Konieczny, S. (2004). Belief base merging as a game. *Journal of Applied Non-Classical Logics*, *14*(3)*,* 275-294.

Konieczny, S., Lang, J., & Marquis, P. (2004). DA2 merging operators. *Artificial Intelligence*, *157*(1-2)*,* 49-79.

Konieczny, S., & Perez, R. P. (2005). Propositional belief base merging or how to merge beliefs/goals coming from several sources and some links with social choice theory. *European Journal of Operational Research*, *160*(3)*,* 785-802.

Kotzinos, D., Pediaditaki, S., Apostolidis, A., Athanasis, N., & Christophides, V. (2005). Online curriculum on the semantic Web: The CSD-UoC portal for peer-to-peer e-learning. In *Proceedings of the 14th International WWW Conference (WWW-05)*, Chiba, Japan.

Krotzsch, M., Vrandecic, D., & Volkel, M. (2005). Wikipedia and the semantic Web – The missing links. In *Proceedings of Wikimania*.

Lamb, B. (2004). Wide open spaces: Wikis, ready or not. *EDUCAUSE Review*, *39*(5)*,* 36-48.

Liebig, T., & Noppens, O. (2005). OntoTrack: A semantic approach for ontology authoring. *Journal of Web Semantics*, *3*(2-3)*,* 116-131.

Magiridou, M., Sahtouris, S., Christophides, V., & Koubarakis, M. (2005). RUL: A declarative update language for RDF. In *Proceedings of the 4th International Semantic Web Conference (ISWC-05)*, Galway, Ireland.

Martin, P., Blumenstein, M., & Deer, A. (2005). Toward cooperatively-built knowledge repositories. In *Proceedings of the 13th International Conference on Conceptual Structures (ICCS-05)*, (LNCS 3596, pp. 411-424). Springer-Verlag

Mathes, A. (2004). Folksonomies-Cooperative classification and communication through shared metadata. *Computer Mediated Communication*, LIS590CMC (Doctoral Seminar), Graduate School of Library and Information Science, University of Illinois Urbana-Champaign.

McGreal, R., & Roberts, T. (2001). *A primer on metadata for learning objects.* Retrieved January 24, 2007, from http://www.elearningmag.com/issues/ct01/earningobjects.asp

Mika, P. (2005). Ontologies are us: A unified model of social networks and semantics. In *Proceedings of the International Semantic Web Conference, (ISWC-05)*, (pp. 522-536).

Miller, E., & Swick, R. (2003). RDF and RDF schema, W3C. Retrieved January 24, 2007, from http://www.w3.org/RDF

Noy, N. F., Kunnatur, S., Klein, M., & Musen, M. A. (2004). Tracking changes during ontology evolution. In *Proceedings of the 3rd International Semantic Web Conference (ISWC-04)*, Hiroshima, Japan.

Ohmukai, I., Hamasaki, M., & Takeda, H. (2005). A proposal of community-based folksonomy with RDF metadata. *Workshop on End User Semantic Web Interaction*, held in conjunction with the 4th International Semantic Web Conference (ISWC-05).

Paavola, S., Lipponen, L., & Hakkarainen, K. (2004). Models of innovative knowledge communities and three metaphors of learning. *Review of Educational Research*, *74*(4), 557-576.

Peters, J. (1995). The Hundred Years War started today: An exploration of electronic peer review. *Internet Research: Electronic Networking Applications and Policy*, *5*(4), 3-9.

Richardson, M., & Domingos, P. (2003). Building large knowledge bases by mass collaboration. In *Proceedings of the 2nd International Conference on Knowledge Capture (K-CAP-03)*, (pp. 129-137).

Schaffert, S., Gruber, A., & Westenthaler, R. (2005). A semantic wiki for collaborative knowledge formation. In *Proceedings of Semantics 2005*, Vienna, Austria.

Slavin, R. E. (1989). Research on cooperative learning: An international perspective. *Scandinavian Journal of Educational Research*, *33*(4), 231-243.

Spyratos, N., & Christophides, V. (2005). Querying with preferences in a digital library. In *Dagstuhl Seminar (no 05182) Federation over the Web*.

Tazzoli, R., Castagna, P., & Campanini, S. (2004). Towards a semantic wiki wiki Web. In *Proceedings of the 3rd International Semantic Web Conference (ISWC-04)*.

Teniente, E., & Olivie, A. (1995). Updating knowledge bases while maintaining their consistency. *VLDB Journal*, *4*(2), 193-241.

Tennison, J., & Shadbolt, N. (1998). APECKS: A tool to support living ontologies. In *Proceedings of the 11th Knowledge Acquisition Workshop (KAW-98)*.

Tzitzikas, Y. (2001). Democratic data fusion for information retrieval mediators. In *ACS/IEEE International Conference on Computer Systems and Applications*, Beirut, Lebanon.

Tzitzikas, Y., & Hainaut, J. L. (2005). How to tame a very large ER diagram (using link analysis and force-directed placement algorithms). In *Proceedings of 24th International Conference on Conceptual Modeling, (ER-05)*, Klagenfurt, Austria.

Tzitzikas, Y., & Hainaut, J. L. (2006). On the visualization of ontologies. In *Proceedings of the 10th International Conference on Advanced Visual Interfaces, (AVI-06)*, Venice, Italy.

Volkel, M., & Oren, E. (2005). *Personal knowledge management with semantic wikis*(Tech. Rep.). AIFB Karlsruhe.

Volkel, M., Winkler, W., Sure, Y., Kruk, S. R., & Synak, M. (2005). SemVersion: A versioning system for RDF and ontologies. In *Proceedings of the 2nd European Semantic Web Conference (ESWC-05)*, Heraklion, Crete.

W3C (2004). *RDF vocabulary description language 1.0: RDF schema*. Retrieved January 24, 2007, from http://www.w3.org/TR/rdf-schema/

Watkins, E. R., & Nicole, D. A. (2006). Named graphs as a mechanism for reasoning about provenance. In *Proceedings of the 8th Asia-Pacific Web Conference*, (APWeb-06), (pp. 943-948), Harbin, China.

Wiederhold, G. (1994). An algebra for ontology composition. In *Proceedings of the 1994 Monterey Workshop on Formal Methods*, (pp. 56-61).

ENDNOTES

1. This fact is based on the successful results of experiments reported in (Gokhale, 1995), where fifty percent of each student's individual grade was based on the average score (of the group members) while the remaining fifty percent of each student's grade was individual.
2. This sketch is by no means a proposed UI design.
3. This option would not make much sense if O is defined by union, but it could be reasonable if O is defined by intersection or quantitatively.
4. Of (or relating to) a synod, where synod is a council or an assembly.
5. The problem of maintaining consistency after updates has been studied in the Database and KR literature (e.g. see (Teniente & Olivie, 1995)) mainly for the single actor case and also in Decision Support Systems (Karacapilidis & Papadias, 2001).
6. This could be one answer to the learning paradox, i.e., to the classical problem of explaining how something new and more complex is created using existing knowledge. Specifically, the composition of inconsistent KBs could yield a new (may bigger and more complex) KB that it is consistent.
7. Related recent projects for building collective KBs include Open Mind (www.openmind.org) and Mind Pixel (www.mindpixel.com).

Yannis Tzitzikas is currently an assistant professor in the Computer Science Dep. at University of Crete (Greece) and associate researcher in Information Systems Lab at FORTH-ICS (Greece). Before joining UofCrete and FORTH-ICS he was postdoctoral fellow at the University of Namur (Belgium) and ERCIM postdoctoral fellow at ISTI-CNR (Pisa, Italy) and at VTT Technical Research Centre of Finland. He conducted his undergraduate and graduate studies (MSc, PhD) in the Computer Science Department at University of Crete. In parallel, he was a member of the Information Systems Lab of FORTH-ICS for about 8 years, where he conducted basic and applied research around semantic-network-based information systems within several EU-founded research projects. His research interests fall in the intersection of the following areas: information systems, information indexing and retrieval, conceptual modeling, knowledge representation and reasoning, and collaborative distributed applications. The results of his research have been published in more than 40 papers in refereed international conferences and journals, and he has received one best paper award.

Vassilis Christophides studied electrical engineering at the National Technical University of Athens (NTUA), Greece. He received his DEA in computer science from the University PARIS VI and his Ph.D. from the Conservatoire National des Arts et Metiers (CNAM) of Paris, France. He is an associate professor at the Department of Computer Science, University of Crete, and affiliated researcher at the Information Systems and Software Technology Laboratory of the Institute of Computer Science Foundation for Research and Technology-Hellas (FORTH-ICS), leading several EU projects related to the integration of heterogeneous and distributed information sources. His main research interests include Semantic Web and peer-to-peer information

management systems, semistructured and XML/RDF data models and query languages as well as description and composition languages for e-services. He has published over 60 articles in international conferences and journals and has served on numerous conferences program committees (ACM SIGMOD, VLDB, EDBT, WWW, ICSW, ICWE, ICWS, ECDL).

Giorgos Flouris is currently a post-doctoral research fellow at the Istituto della Scienza e delle Tecnologie della Informazione (ISTI) of CNR in Italy, under an ERCIM "Alain Bensoussan" post-doctoral fellowship. He holds a BSc in mathematics from the University of Athens and an MSc and a PhD in computer science from the University of Crete. His research interests lie in the areas of knowledge representation and reasoning, belief revision, the semantic web and ontology evolution. Dr. Flouris has published about 20 research papers in peer-reviewed workshops, conferences and journals and has received a number of scholarships and awards, including a Best Paper Award in STAIRS-06. He is currently involved in the EU projects KP-Lab and CASPAR and has served as an external PC member in several workshops, conferences and journals; moreover, he has organized the IWOD-07 workshop, a part of ESWC-07 conference.

Dimitris Kotzinos is an adjunct assistant professor at the Computer Science Department at the University of Crete and an affiliated researcher at the Information Systems and Software Technology Laboratory of the Institute of Computer Science Foundation for Research and Technology-Hellas (FORTH-ICS). He holds a PhD on the topic of "Application of digital map technologies on developing internet based Advanced Traveler Information Systems (ATIS)" from the Department of Production and Management, Technical University of Crete, Greece (2001) and a MSc in transportation, from the Civil Engineering Department, Ohio-State University, Columbus, USA (1996). His BS is in computer science, Department of Computer Science, University of Crete, Greece. (1994). His main research interests include development of methodologies, algorithms and tools for Web based information systems, portals and Web services. Especially applications of the above in the fields of: e-learning, geographic information portals, real-time advanced traveler information systems (A.T.I.S.). He has published over 25 paper in various journals, conferences and workshops and serves as a program committee member and reviewer for various conferences and journals.

Hannu Markkanen, MSc (Eng.), is a principal research lecturer in media engineering at EVTEK University of Applied Sciences, Espoo, Finland. Currently he is the technical coordinator of the KP-Lab project of the EU's Information Society Technologies -program. During the last 15 years he has coordinated international R&D projects in the field of learning technologies and has published over 20 papers and articles in professional conferences and publications. His current R&D interests are in the area of semantic learning and collaboration environments.

Dimitris Plexousakis is an associate professor and Chair at the Department of Computer Science, University of Crete and a researcher at the Information Systems Laboratory of the Institute of Computer Science, FORTH in Greece. His research interests span the following areas: knowledge representation, knowledge base design; distributed database systems and databases on the Web; the Semantic Web; formal reasoning systems, applications of artificial intelligence in database systems; business process and e-service modeling. He has published over 60 articles in international conferences and journals and has served on the program committees of numerous international conferences and journal editorial boards. He was the executive chair of the 9th International Conference on Extending Database Technology (March 14-18, 2004, Heraklion, Greece) and the Program Committee co-Chair of the 3rd International Semantic

Web Conference (November 7-11, 2004, Hiroshima, Japan). He is leading the ERCIM Working Group on the Semantic Web. He is on the editorial board of the Journal of Web Semantics *and has co-chaired a number of workshop on subjects related to Semantic Web research, including the Workshop on E-Services and the Semantic Web and the Workshop on Data Integration and the Semantic Web.*

Nicolas Spyratos is full professor at the Department of Computer Science, University of Paris 11 (Orsay Center), and head of the database group. Prior to joining the University of Paris, Prof. Spyratos worked as a research scientist for the CNRS and INRIA, in France, Bell Northern Research in Canada (now, Northern Telecom), and IBM in Greece. His current research is in the area of databases and conceptual modelling, and more recently in the area of digital libraries. He is the author of over 100 scholarly publications, and has participated in several European Union projects.