

Emergent Knowledge Artifacts for Supporting Trialogical E-Learning

Yannis Tzitzikas^{1,2}, Vassilis Christophides^{1,2}, Giorgos Flouris², Dimitris Kotzinos^{1,2}, Hannu Markkanen³, Dimitris Plexousakis^{1,2}, Nicolas Spyratos⁴

¹ Computer Science Department, University of Crete, GREECE

² Institute of Computer Science, FORTH-ICS, GREECE

³ EVTEK University of Applied Sciences, FINLAND

⁴ Laboratoire de Recherche en Informatique, Université de Paris-Sud, FRANCE

Emails: {tzitzik | christop | fgeo | kotzino | dp}@ics.forth.gr,
hannu.markkanen@evtek.fi, spyratos@lri.fr

Abstract. This paper 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 the paper analyzes and discusses related techniques and identifies issues that are worth further research.

1 Introduction

Classical learning theories are based either on the *knowledge acquisition* metaphor (i.e., a learner individually internalizes a body of knowledge) or on the *social participation* metaphor (i.e., 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, i.e., 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*) focus on the social processes by which learners collectively enrich/transform their individual and shared cognition. According to TL, knowledge creation activities rely heavily on the use, manipulation and evolution of shared knowledge artifacts externalizing a body of (tacit or explicit) knowledge [29]. 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 and thus,

finally reach a common understanding of the problem at hand. We could therefore consider as cornerstone of trialogical learning the notion of shared objects of activity, a notion that is quite general to accommodate the requirements of various application contexts. For instance, a video that records how group members carry out their tasks, could be considered as a shared knowledge artifact which 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 *knowledge artifact* to refer to what is being created and/or shared by a group of learners (and 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 [9].

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 world views are still in their infancy. In this paper, we investigate various ways to build emerging knowledge spaces. We have used the trialogical learning paradigm for eliciting the functional requirements. In particular, we focus on the various methods 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: Section 2 describes a TL scenario for collaborative knowledge creation, and Section 3 discusses the underlying principles and interactions. Section 4 describes various ways to build emerging knowledge artifacts from individual group knowledge (of various forms), and identifies knowledge management requirements. Finally, Section 5 summarizes and concludes the paper.

2 Motivating Scenario for Trialogical Learning

2.1 Collaborative Literature Review and Annotation

A set of N research papers, say $P = \{p_1, \dots, p_N\}$, is given to a set of K learners $A = \{a_1 \dots a_K\}$ who could be students, researchers, or co-workers 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, i.e. 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 a MSc student should do in order to write the state-of-the-art of his MSc thesis. Of course, this scenario does not fall into dialogical learning, but is rather an instance of monological learning (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 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 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.

2.2 Grading and Progress Assessment of Individuals and Groups

A related rising question is whether the "quality" of the result of this collaboration (i.e. of O and $d(p)$'s) should be measured and if yes how. We can identify two broad cases. According to the first, there is an external (human or machine) observer who can grade the result, while according to the second there is not any external party. For instance, we may assume that there is a certain "solution" ontology (ideal or criterion), denoted $O^{(i)}$ that is unknown for members of the group. For example, $O^{(i)}$ could have been provided by a tutor if there is one (or the tutor might have provided a set of admissible ontologies instead of one ontology). Subsequently, appropriate metrics could be employed in order to measure the "distance" between $O^{(i)}$ and O_{s_i} and at every point in time (state s_i), so that the members of the group can judge if they progress or not. Of course not only the group work but also the individual work could be graded. Recall that according to [30, 13], for effective collaborative learning, there must be "group goals" and "individual accountability"¹.

In the case where there is not any external party we could probably only measure the degree of agreement between the members of the group. If O_A expresses the knowledge that all members of A accept to be correct, then the bigger O_A is, the better the group goes (assuming there is not any other constraint like C in the previous scenario).

3 Emergent Knowledge Artifacts Spaces

This section discusses issues that are important for supporting the previous scenario. In particular, Section 3.1 introduces personal and shared knowledge

¹ Based on the successful results of experiments reported in [13]: 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.

artifacts and clarifies their relation, while Section 3.2 shows how a set of learners can interact on the basis of their personal and shared knowledge artifacts. It also discusses synoptically additional issues.

3.1 Personal versus 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 knowledge base of an actor a . Now let KB_A denote the "shared" (or common) knowledge base of a set of actors A . The important issue here is the relation between KB_A and KB_a (for $a \in A$). Below we identify three broad cases:

- *UNION-case*. Here KB_A is obtained by taking the union of the KBs of all participants, i.e.: $KB_A = \cup\{KB_a \mid a \in A\}$. Note that 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 a very strict 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, i.e.: $KB_A = \cap\{KB_a \mid a \in A\}$, so it comprises statements "accepted" by every participant.
- *QUANTITATIVE-case*. Here KB_A is defined by a quantitative method, e.g. it may comprise all sentences that are accepted by at least a percentage of the actors. Obviously, UNION and INTERSECTION are special cases of this case.

3.2 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) and supports personal and shared knowledge artifacts. Figure 1 sketches a possible UI² for that 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. In the left area each learner is free to do whatever she wants, so everything is editable in that area. The right area shows the shared artifacts and this area is the key point for collaboration and for supporting trialogical e-learning. For instance, and assuming the scenario described earlier, each user may develop her own ontology at the left area, while the right window shows the group ontology O (according to the method that O is derived from the personal ontologies).

² This sketch is by no means a proposed UI design.




Fig. 1. An indicative UI for trialogical E-learning

The relationship between personal space and group space is very important. The button labeled by "→" allows a user to copy the desired parts from 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 is immediately reflected (propagated) to the group ontology (e.g. blackboard-based collaboration). Symmetrically, an option that keeps the button "←" permanently pressed would propagate the changes on O to the personal space³. Deletions are handled analogously and are discussed in Section 4.2. We could call systems (and UIs) that allow this kind of collaboration/interaction *synodic*⁴.

Above we have sketched the basics of a trialogical e-learning scenario. Of course, the scenario (and the UI) can be enriched with a plethora of auxiliary functionalities. Below we identify the most important ones according to our opinion:

- The *group space* view could be *customizable*, e.g. 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. In general, the shared knowledge base could be defined with a set theoretic expression over subsets 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. Moreover, the group space could be *optionally managed* by a person whose role would be to accept or reject the changes that the participants forward to the group ontology.
- The *provenance* of every statement should be saved and be available at any time (e.g. this link was added by learner a_2). Moreover, the participants should be able to *annotate* every element of their personal or group space. The 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 (or what percent of the actors agree with this). More-

³ This is not reasonable 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.

- over the *visualization* of knowledge artifacts is a very important, challenging and open issue (some related issues are discussed in brief in [24, 35]).
- The UI could be enriched with *teleconferencing* services allowing the participants to discuss in real-time while using the system.

4 Synthesizing Knowledge Bases

To support the scenario described in Section 2, we need to support the formation and evolution of A , of P , of O , and $d(p)$'s. In order to identify the distinctive knowledge management requirements for supporting e-trialogical learning, we will first present an approach for supporting personal and shared knowledge artifacts and then we will investigate various forms of knowledge bases starting from the very simple ones. The reason for trying to identify the key knowledge management requirements (that originate from TL), is to investigate how we could support them by extending accordingly the core knowledge management technologies (and not by developing yet another e-learning application).

4.1 Supporting Personal and Shared Knowledge Artifacts

Now we will divide the personal space of an actor into two spaces: one private and one public. The group (shared) space is derived from the public personal spaces of the actors.

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" (e.g. construct or update operation) concerning his personal space. The second, denoted by a_i , is associated to every statement he has forwarded to the group space. Let KB_i^p denote the knowledge base comprising all statements with identifier a_i^p , and KB_i denote the knowledge base of statements with identifier a_i . Normally, it should be $KB_i \subseteq KB_i^p$, that is the public personal base of a user should be subset of the personal private base of that user. However, in social life sometimes persons forejudge or "pretend" that they accept facts although they don't really believe them (e.g. because all other persons do, or for strategic reasons). In such cases the relationship $KB_i \subseteq KB_i^p$ does not hold. For this reason, and in order to leave learners free, we shouldn't impose any constraint among KB_i and KB_i^p .

The important point here is that the synthesis (or amalgamation) of all KB_i 's forms the shared artifacts of the group (i.e. the shared artifacts according to trialogical learning). Let's now return to our application scenario, and suppose the case where there is one tutor who has also provided to the learners a preliminary version of the ontology O_{pre} (on which the learners should work on). We could capture this case by considering that initially it holds $KB_i^p = O_{pre}$ for each $i = 1..K$.

4.2 KB = A Set of Words

In order to identify the distinctive knowledge management requirements for supporting trialogical learning (if any), we will start from very simple forms of knowledge bases.

Suppose that a knowledge base 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 every actor $a_i \in A$ we have two knowledge bases: KB_i^p and KB_i . The first is a set of pairs of the form (w, a_i^p) while 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:

1. 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 himself the word w to his private base.
2. The private base of the user is updated accordingly.
3. Probably (or optionally) a pair (w, a_i) should be created.

Let's now suppose that the user *deletes* one element w of his private knowledge base. If the user had "published" w in the past, i.e. if a pair (w, a_i) exists, then the system should ask the user if the pair (w, a_i) should be deleted or not. This case suggests that it would be more informative if the UI for each actor a_i were divided into 3 areas: one showing KB_i^p , one KB_i , and one for KB_A , as it is depicted in Figure 2. This would allow monitoring and controlling the contents of KB_i .




Fig. 2. An indicative UI for triological E-learning

Let's now investigate how the "shared" knowledge base might be defined. Let KB_A denote the KB obtained by taking the union of the public bases of all actors, i.e. $KB_A = \bigcup_1^K KB_i$. 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 at run-time the group KB by various methods (by union, intersection or any other) as shown below.

- The UNION case comprises all words w such that $|for(w)| \geq 1$, specifically:

$$KB_{\cup A} = \{w \mid for(w) \subseteq A\}$$

- The INTERSECTION-case comprises all words w such that $|for(w)| = K$, specifically:

$$KB_{\cap A} = \{w \mid for(w) \supseteq A\}$$

- The z -PERCENT case comprises all words w such that $|for(w)|/K \geq z$, specifically:

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

- The case where a user wants to see the group ontology as derived by considering only a subset A' of A can be captured by the above 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 knowledge base" we want. So such a representation could be adopted for the physical layer of the repository.

Grading (assessing progress)

Let W and W' be the set of words stored in two knowledge bases KB and KB' respectively. We can define the distance between two knowledge bases KB and KB' on the basis of W and W' . For instance, we can use the symmetric difference, i.e. $dist(KB, KB') = |W \setminus W'| + |W' \setminus W|$, the Dice coefficient, i.e. $dist(KB, KB') = 1 - \frac{|W \cap W'|}{|W \cup W'|}$, or any other metric.

4.3 KB = A Binary Relation

Now suppose that a KB is a binary relation R over a set of elements T i.e. $R \subseteq T^2$. Let r denote an element of a R , e.g. $r = (t, t')$ where $t, t' \in T$. In our application scenario, this corresponds to the case where the ontology (that the learners have to create) is a graph of keywords.

We can define the personal and group knowledge bases as we did earlier (e.g. $KB_{\cup A} = \{r \mid for(r) \subseteq A\}$). 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 R_i and T_i (of course $R_i \subseteq T_i^2$). It follows that we can define shared knowledge bases (e.g. $KB_{\cup A}$ and $KB_{\cap A}$) not only for R but also for T .

4.4 KB = A Binary Relation with Second Order Properties

Here we consider the case where a KB is a binary relation R over a set of elements T ($R \subseteq T^2$), with the extra rule or constraint that this relation satisfies a property (e.g. that R is reflexive, symmetric, antisymmetric, transitive, etc). These extra properties can be seen as derivation rules (inferences) or constraints. For instance, note that the case where R is a preorder (i.e. a reflexive and transitive relation) captures the case of taxonomies. So in our application scenario, this corresponds to the case where the ontology (that the learners have to create) has the form of a taxonomy. Supporting this scenario is actually supporting collaborative (and triological) taxonomy construction.

We could model inferences (e.g. transitivity) as follows. We can consider a KB as a set of sentences S and we make the assumption that there is a consequence operation $Cons$ that models inference services ($S \subseteq Cons(S)$). Also note that axioms could be modeled by the notion of consistency.

It follows that for each $i = 1..K$ we have KB_i , KB_i^p , $Cons(KB_i)$ and $Cons(KB_i^p)$. A "shared" knowledge base can be defined on the basis of KB_i or on the basis of $Cons(KB_i)$. The resulting shared knowledge base can be different in each case, as shown in the example of Figure 3 where $KB_{\cap\{1,2,*\}}$ has been used to denote that $Cons(KB_1)$ and $Cons(KB_2)$ were used for the definition of $KB_{\cap\{1,2\}}$.




Fig. 3. Local Reasoning and Group KBs

Total Order Consider now the case where R is a total order. For instance, consider the case where learners have to rank a set of available options T in order to come up with some decision. For example, the learners may have to rank a set of keywords or a set of papers according to their significance or importance. In addition, suppose a questionnaire comprising multiple choice questions where more than one choices are correct for each question but the tutor asked from the group to mark only one choice (the most appropriate). Also notice that the case of total orders captures the selection process of peer-reviewed scientific conferences and journals. Here the shared (group) knowledge base can be obtained by aggregating the "rankings" of the learners. For doing an aggregation of this kind, we could adopt various techniques (mainly coming from the area of Social Choice), like plurality ranking, Borda [7] ranking, Condorcet [8]

ranking or Kemeny Optimal Aggregation [20], but we shouldn't forget the Arrow's impossibility theorem [2]. A Borda-like technique for aggregating weakly ordered subsets of a set which could be used for our purposes, is described in [33]. 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 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 $\langle o_1, o_2, o_3 \rangle$ meaning that o_1 is preferable to o_2 which is preferable to o_3 , we may have input of the form $\{score(o_1) = 5, score(o_2) = 3\}$).

In the above scenario the set T is not part of the created knowledge (in other words, it preexists). A scenario where T does not preexist but is rather part of the created knowledge follows. 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 TEL-CoPs'06. Each one proposes a structure, i.e. 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, i.e. with one structure either accepted by all of them or by most of them. As it wouldn't 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 to a structure that is probably better than what each one could do by himself (of course aposties may occur). An alternative method to support this scenario follows. Suppose that the paper to be submitted should have exactly 7 sections. Let T be the pieces of texts that all actors have provided (i.e. $T = \cup_1^K T_i$), e.g. 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 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.

4.5 KB = An RDF-based Repository

Suppose now the case that the learners have to create an ontology-based repository (ontology plus ontology-based metadata). A repository of this kind has the form of a conceptual graph. According to RDF [27, 4], this graph can be seen as a set of RDF triples which actually defines a directed graph consisting of 3 kinds of relations (instanceOf, isA and property). So we could write

$KB = (R_{in}, R_{isa}, R_p)$, where R_{in} comprises instanceOf relationships, R_{isa} comprises isa relationships, and R_p comprises property relationships. Note that the *isa* relation (R_{isa}) models a transitive relation so the issues discussed in Section 4.4 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 [17].

Notions of consistency could arise in such a setting. If inconsistency arises in one individual (personal) KB, then the user is responsible for making what is necessary for reaching a consistent one⁵. However, one can easily see that although each individual personal KB may be "consistent", the group ontology may be not. Who and how should react in that case? Should the system allow such cases and if yes is there anything it could do for aiding actors to overcome this problem? One first remark is that it wouldn't be flexible to forbid inconsistent group KBs. So the system should allow inconsistent group KBs but it should be at least able to detect inconsistencies and indicate them to the actors. If we allow inconsistency also in the personal KBs, then another interesting case may occur: the individual KBs could be inconsistent while the group KB is consistent⁶.

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's 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, then the above definitions of group KBs coincide with the original ones.

Similarly, we could 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 in the learner's mental state 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, as mentioned in Section 4.4. Furthermore, it allows the adaptation of works related to belief merging [21], [23], [22] in our aggregation context, by facilitating the formal description of notions like "weakening", "conceding" and "negotiating" [21], the development of *arbitration* or *majority* merging operators [23] and the definition of distances and aggregation functions [22].

⁵ The problem of maintaining consistency after updates have been studied in the Database & KR literature (e.g. see [32]) but mainly for the single actor case.

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.

Note that unlike traditional approaches conceiving ontologies as thorough engineering artifacts issued by strict design process and policies, in TL ontology creation and evolution can be seen as a social process where learners collectively improve their individual and shared understanding through social interaction. 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 [1]). Ontologies would thus become an emergent effect of open-ended interactions within or across groups of individuals as opposed to be a firm commitment of a small group of domain experts (for more see [26]).

Further Issues As we step up the expressive power of the representation framework additional issues arise:

- For instance, knowledge change and evolution raises various issues e.g. the distinction between *update* and *revision* (in the sense defined in [19]), as well as the applicability of belief revision theories to ontology evolution (e.g. see [11]).
- Measuring the distance between two knowledge bases (e.g. for grading as described in Section 2) may not be enough. It will be also important (e.g. for learning purposes) to compute and show the difference, or *delta*, between two knowledge bases. Some approaches for computing deltas of RDF graphs are described in SemVersion [36], PromptDiff [28] and [3].

Furthermore, as the number of actors scales up, additional issues arise, e.g. 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, especially because no one ever tried to impose a structure or any form of control on that. It follows that link analysis techniques (either applied on social networks, or on articulated knowledge bases [5, 15], or on large knowledge bases [34]) are also expected to be useful in large-scale collaborative knowledge creation. The provision of personalized services is also very useful in large-sized knowledge bases [31].

As a final remark, note that the need for defining separate knowledge spaces and for combining them has been identified in several contexts also in 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 [38, 10, 16, 14]), and the more recent one is that of *named graphs* [6, 37]. In this paper we go one step further and we stress the need for synthesizing such knowledge spaces.

At last, we should remark that workflow issues are orthogonal to the issues we discussed so far. The issues we elaborated so far are raised in almost every step of a workflow process if that step should be carried out collaboratively.

5 Epilogue

This paper described a specific scenario for collaborative knowledge creation in the spirit of the triological learning paradigm. According to this scenario the group knowledge base is formed by combining the KBs 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. 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 we 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.

Summarizing, triological 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 mainly for the case where there is a commonly accepted conceptualization and world view. Methodologies and technologies that allow diversification and flexible amalgamation of different world views have not emerged so far. Areas of knowledge management that are related (in principle) to triological 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 (co-funded by the IST programme of the EU 6). The implementation will be based on Semantic Web technologies specifically on the RDF Suite [12, 18, 25].

Acknowledgements

This research was conducted within the Knowledge Practices Laboratory (KP-Lab) project co-funded by the IST programme of the EU 6 R&D Framework programme.

References

1. K. Aberer, T. Catarci, P. Cudré-Mauroux, T. S. Dillon, S. Grimm, M. Hacid, A. Illarramendi, M. Jarrar, V. Kashyap, M. Mecella, E. Mena, E. J. Neuhold, A. M. Ouksel, T. Risse, M. Scannapieco, F. Saltor, L. De Santis, S. Spaccapietra, S. Staab, R. Studer, and O. De Troyer. "Emergent Semantics Systems". In *Procs. of the 1st Intern. IFIP Conference on Semantics of a Networked World (ICSNW 2004)*, pages 14–43, 2004.
2. Kenneth J. Arrow. *Social Choice and Individual Values*. 1951.
3. Tim Beners-Lee and Connolly. "Delta: An Ontology for the Distribution of Differences Between RDF Graphs", 2004. <http://www.w3.org/DesignIssues/Diff> (version: 2004-05-01).

4. Dan Brickley and R. V. Guha. "Resource Description Framework (RDF) Schema specification: Proposed Recommendation, W3C", March 1999. <http://www.w3.org/TR/1999/PR-rdf-schema-19990303>.
5. Sergey Brin and Lawrence Page. "The Anatomy of a Large-scale Hypertextual Web Search Engine". In *Proceedings of the 7th International WWW Conference*, Brisbane, Australia, April 1998.
6. J. Carroll, C. Bizer, P. Hayes, and P. Stickler. Named graphs, provenance and trust. *Proceedings of the 14th international conference on World Wide Web*, pages 613–622, 2005.
7. JeanCharles de Borda. "Memoire sur les Elections au Scrutin", 1781. Histoire de l'Academie Royale des Sciences, Paris.
8. Marquis de Condorcet. "Essai sur l'application de l'analyse a la probabilite des decisions rendues a la pluralite des voix", 1785. Paris.
9. John Domingue, Enrico Motta, Simon Buckingham Shum, Maria Vargas-Vera, Yannis Kalfoglou, and Nick Farnes. "Supporting Ontology driven Document Enrichment within Communities of Practice". In *Procs of the 1st intern. conference on Knowledge Capture, K-CAP'01*, pages 30–37, New York, NY, USA, 2001. ACM Press.
10. J. Euzenat. "Corporate Memory through Cooperative Creation of Knowledge based and Hyper-documents". In *Proceedings of the 10th Knowledge Acquisition Workshop, KAW'96*, November 1996.
11. Giorgos Flouris. "On Belief Change and Ontology Evolution". PhD thesis, Computer Science Department, University of Crete, Greece, 2006.
12. FORTH-ICS. The ics-forth rdfsuite: High-level scalable tools for the semantic web, 2005. <http://139.91.183.30:9090/RDF/>.
13. Anuradha A. Gokhale. "Collaborative Learning Enhances Critical Thinking". *Journal of Technology Education*, 7(1), 1995.
14. R. Guha, R. McCool, and R. Fikes. Contexts for the Semantic Web. *Proceedings of the 3rd International Semantic Web Conference, Hiroshima*, 2004.
15. Lin Guo, Feng Shao, Chavdar Botev, and Jayavel Shanmugasundaram. "XRANK: Ranked Keyword Search over XML Documents". In *Int. Conf. on Management of Data, SIGMOD'2003*, San Diego, California, June 2003.
16. J. Jannink, S. Pichai, D. Verheijen, and G. Wiederhold. "Encapsulation and composition of ontologies". In *Proceedings of 1998 AAAI Workshop on AI & Information Integration*, 1998.
17. Zoi Kaoudi, Theodore Dalamagas, and Timos Sellis. "RDFSculpt: Managing RDF Schemas Under Set-Like Semantics". In *Procs of the 2nd European Semantic Web Conference 2005 (ESWC05)*, Crete, Greece, 2005.
18. G. Karvounarakis, V. Christophides, and D. Plexousakis. "RQL: A Declarative Query Language for RDF". In *Eleventh International World Wide Web Conference (WWW)*, Hawaii, USA, May 2002.
19. H. Katsuno and A. O. Mendelzon. "On the Difference between Updating a Knowledge Base and Revising it". In *Proceedings KR-91*, pages 380–395, 1991.
20. John Kemeny. "Mathematics without Numbers". *Daedalus*, 88:571–591, 1959.
21. Sebastien Konieczny. Belief base merging as a game. *Journal of Applied Non-Classical Logics*, 14(3):275–294, 2004.
22. Sebastien Konieczny, Jerome Lang, and Pierre Marquis. Da2 merging operators. *Artificial Intelligence*, 157(1-2):49–79, 2004.
23. Sebastien Konieczny and Ramon Pino Perez. 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, 2005.

24. Thorsten Liebig and Olaf Noppens. "OntoTrack: A Semantic Approach for Ontology Authoring". *Journal of Web Semantics*, 3(2-3):116–131, 2005.
25. M. Magiridou, S. Sahtouris, V. Christophides, and M. Koubarakis. "RUL: A Declarative Update Language for RDF". In *Procs. 4th Intern. Conf. on the Semantic Web (ISWC-2005)*, Galway, Ireland, November 2005.
26. Peter Mika. "Ontologies Are Us: A Unified Model of Social Networks and Semantics". In *Procs of the Intern. Semantic Web Conference, 2005*, pages 522–536, 2005.
27. E. Miller, R. Swick, and D. Brickley (editors). *RDF and RDF Schema, W3C*, 2003. <http://www.w3.org/RDF>.
28. Natalya F. Noy, Sandhya Kunnatur, Michel Klein, and Mark A. Musen. "Tracking Changes During Ontology Evolution". In *Third International Conference on the Semantic Web (ISWC-2004)*, Hisroshima, Japan, 2004.
29. S. Paavola, L. Lipponen, and K. Hakkarainen. "Models of Innovative Knowledge Communities and Three Metaphors of Learning". *Review of Educational Research*, 74(4):557–576, 2004.
30. R. E. Slavin. "Research on Cooperative Learning: An International Perspective". *Scandinavian Journal of Educational Research*, 33(4):231–243, 1989.
31. N. Spyrtatos and V. Christophides. "Querying with Preferences in a Digital Library". In *Dagstuhl Seminar (No 05182) Federation over the Web*, May 2005.
32. Ernest Teniente and Antoni Olivie. "Updating Knowledge Bases while Maintaining their Consistency". *VLDB Journal*, 4(2):193–241, December 1995.
33. Yannis Tzitzikas. "Democratic Data Fusion for Information Retrieval Mediators". In *ACS/IEEE International Conference on Computer Systems and Applications*, Beirut, Lebanon, June 2001.
34. Yannis Tzitzikas and Jean-Luc Hainaut. "How to Tame a Very Large ER Diagram (using Link Analysis and Force-Directed Placement Algorithms)". In *Proceedings of 24th Int. Conf. on Conceptual Modeling, ER'2005*, Klagenfurt, Austria, October 2005.
35. Yannis Tzitzikas and Jean-Luc Hainaut. "On the Visualization of Ontologies". In *Proceedings of the X Int. Conf. on Advanced Visual Interfaces, AVI'2006*, Venice, Italy, May 2006.
36. Max Volkel, Wolf Winkler, York Sure, Sebastian Ryszard Kruk, and Marcin Synak. "SemVersion: A Versioning System for RDF and Ontologies". Heraklion, Crete, May 29 – June 1 2005. *Procs. of the 2nd European Semantic Web Conference, ESWC'05*.
37. E. R. Watkins and D. A. Nicole. Named graphs as a mechanism for reasoning about provenance. In *Proceedings of the 8th Asia-Pacific Web Conference, APWeb'2006*, pages 943–948, Harbin, China, 2006.
38. Gio Wiederhold. "An Algebra for Ontology Composition". In *Proceedings of 1994 Monterey Workshop on Formal Methods*, pages 56–61, September 1994.