Ontology Dynamics with Multimedia Information: The BOEMIE Evolution Methodology *

Silvana Castano[1], Sofia Espinosa[2], Alfio Ferrara[1], Vangelis Karkaletsis[3], Atila Kaya[2], Sylvia Melzer[2], Ralf Möller[2], Stefano Montanelli[1], Georgios Petasis[3]

[1] Università degli Studi di Milano,
DICo, 10235 Milano, Italy,
{castano, ferrara, montanelli}@dico.unimi.it

[2] Hamburg University of Technology,
Institute of Software Systems,
21079 Hamburg, Germany,
{sofia.espinosa,at.kaya,sylvia.melzer,r.f.moeller}@tuhh.de

[3] Software and Knowledge Engineering Laboratory,
Institute of Informatics and Telecommunications,
National Centre for Scientific Research (N.C.S.R.) "Demokritos",
P.O. BOX 60228, Aghia Paraskevi, GR-153 10, Athens, Greece.
{vangelis, petasis}@iit.demokritos.gr

Abstract. In this paper, we present the ontology evolution methodology developed in the context of the BOEMIE [1] project. Ontology evolution in BOEMIE relies on the results obtained through reasoning for the interpretation of multimedia resources in order to evolve (enhance) the ontology, through population of the ontology with new instances, or through enrichment of the ontology with new concepts and new semantic relations.

1 Introduction

Ontology learning is a wide domain of research, involving methods and techniques for the acquisition of an ontology from semantic information/conceptual knowledge extracted from a domain. Being closely related to the field of knowledge acquisition, a significant amount of the work has been presented in the bibliography that concentrates on the task of knowledge acquisition from text, through the re-use of widely adopted natural language processing and machine learning techniques [1, 2]. The methodology proposed by the BOEMIE project tries to extend existing approaches by considering modalities beyond text, such as still image, video and audio. The proposed methodology can be separated into three major components:

– *A multimodal information extraction engine.* This information extraction engine is responsible for extracting instances of "primitive" concepts as well as instances of relations between concept instances, from documents belonging to the textual, still

* This paper has been partially funded by the BOEMIE Project, FP6-027538, 6th EU Framework Programme.

[1] http://www.boemie.org

image, video and audio modalities. “Primitive” concepts (known as “mid-level” concepts within BOEMIE) are concepts whose instances can be directly identified in corpora of a specific modality. For example, in the textual modality the name or the age of a person is a mid-level (or “primitive”) concept, as instances of these concepts can be associated with relevant text portions. On the other hand, the concept person is not a mid-level concept, as it is a “compound” concept that has other concepts as properties, like the person name, age, gender, etc. “Compound” concepts are referred as high-level concepts within the BOEMIE project, and cannot be directly identified in a corpus and thus associated with a corpus segment.

- *A semantic interpretation engine*, responsible for producing one or more explanations of an event described in a document. Semantic interpretation operates on the instances of mid-level concepts and relations between them extracted by the information extraction engine, in order to create instances of high-level concepts that explain the event, always according to the domain ontology. Semantic interpretation is performed through standard (i.e. induction) but also non-standard (i.e. abduction) reasoning services and is formalised as a two-level process. During the first level, semantic interpretation is performed on the extracted information (mid-level concept instances/relations) from a single modality only, in order to form modality specific high-level concept instances. At a second level, the modality specific high-level instances are fused, in order to produce high level concept instances that are not modality specific, and contain information extracted from all involved modalities.
- *An ontology evolution toolkit*, which uses the results obtained through reasoning at the interpretation phase in order to evolve (enhance) the ontology, through population of the ontology with instances, or through enrichment of the ontology with new concepts and new relation types.

In this paper, we focus on the BOEMIE evolution methodology and the role of semantic interpretation for ontology evolution. Then, we describe how ontology population and enrichment are articulated, by discussing the main design principles and by describing samples evolution scenarios. The paper is organized as follows: a description of the ontology evolution methodology in BOEMIE is given in Section 2. Methods and techniques for multimedia interpretation by reasoning are discussed in Section 3. The evolution activities of population and enrichment are presented in Sections 4 and 5, respectively. In Section 6, we discuss the original contributions of the present work. Finally, we give our concluding remarks in Section 7.

2 Ontology Evolution in BOEMIE

Ontology evolution within BOEMIE uses as input the results of the semantic representation performed upon information extracted from multiple modalities (combined through fusion). In order to be able to deal with all possible situations requiring evolution, a pattern-driven approach is adopted. Typical input to the evolution toolkit is in the form of ABoxes, containing the results of the semantic representation of the fused extracted information. These results typically include instances of mid-level concepts, relation instances between mid-level concept instances, high-level instances, relation instances between instances of high-level concepts, and possibly instances of the “unknown” mid-level concept. According to the information contained in ABoxes constituting the input, the evolution pattern selector selects the most prominent evolution pattern, triggering either ontology population or ontology enrichment.

The BOEMIE ontology evolution methodology is shown in Figure 1.

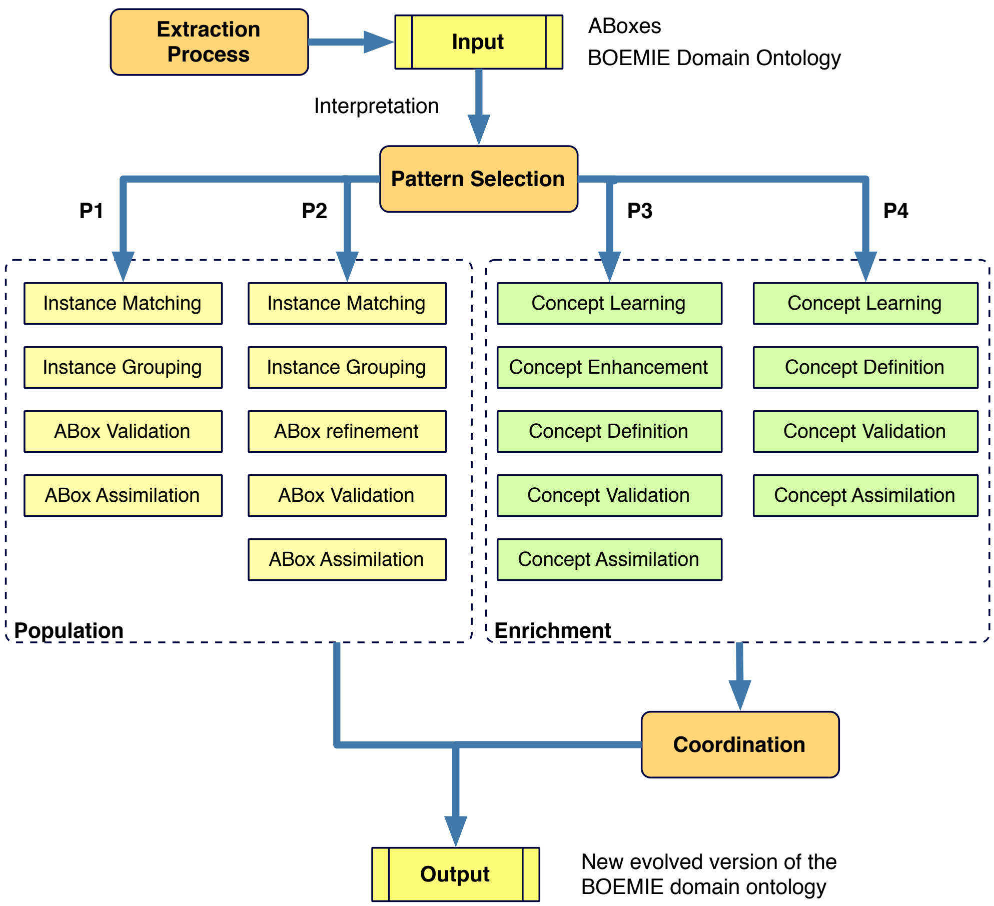

Fig. 1. The BOEMIE evolution methodology

Ontology population is the activity of adding new instances into the ontology and it is performed every time at least one explanation can be found for a multimedia resource.

Ontology enrichment is the activity of extending the ontology, through the addition of new elements (e.g., concepts, relations, properties). Ontology enrichment is performed every time the background knowledge is not sufficient to explain the extracted information from the processed multimedia documents. Thus, the ontology enrichment activity is expected to extend this background knowledge through the addition of new ontology elements.

Coordination is the activity of producing a log of the changes introduced into the new evolved version of the ontology with respect to the initial version and of defining and updating the mapping knowledge between the domain ontology and other related external knowledge sources supporting the enrichment. Since this activity is affected by the changes of the background knowledge, it is executed after the enrichment activity.

In the remainder of the paper, we focus on the reasoning activity for resource interpretation and on the activities of ontology population and enrichment, by providing some relevant example.

2.1 Evolution Patterns

The BOEMIE evolution approach is featured by two main requirements:

- the capability of classifying the different situations that trigger ontology evolution by characterizing the results of the semantic interpretation process (i.e., the information specified in the incoming A-Box) with respect to the background knowledge;
- the definition of an appropriate activity articulation to correctly modify the ontology in each specific evolution situation.

The expected result of the semantic interpretation is a single explanation for a multimedia resource, that is, the resource is instance of only one high-level concept. However, other situations can occur when the background knowledge is not sufficient, characterized either by more than one possible explanation for the same multimedia resource or by the absence of explanations, meaning that no high-level concept can be found in the ontology for describing such a resource. Finally, we can also have the case where not only the high-level concept describing the resource is missing, but also for one or more elements identified in the resource a mid-level concept can not be assigned.

To take these requirements into account, four different *evolution patterns* (see Figure 2) have been identified for ontology evolution in BOEMIE. An evolution pattern determines the characteristics of the input ABox it deals with, defines the kind of evolution to be performed over the ontology (i.e., population or enrichment), and is articulated into a set of activities for implementing all the required changes. Population patterns (P1 and P2) describe the situation where the interpretation has found one or more high-level concepts explaining the resource and, thus, the ontology population activity is performed. Enrichment patterns (P3 and P4) describe the situation where no high-level concepts explaining the resource are found in the ontology, thus triggering ontology enrichment to acquire this missing knowledge. Pattern P4 has been conceived to deal with situations where not only the high-level concept is missing (like

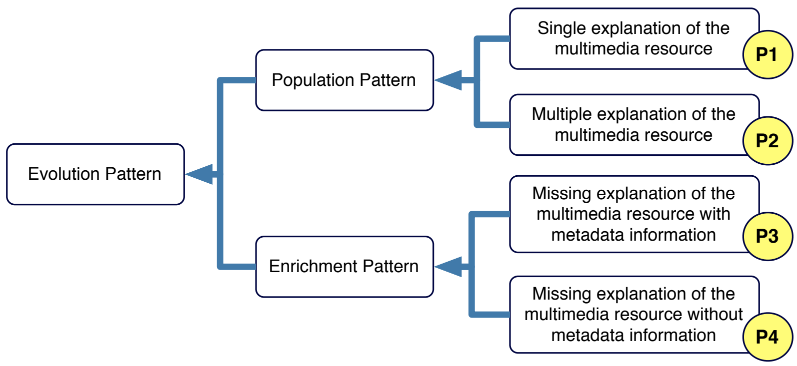

Fig. 2. BOEMIE evolution patterns

P3) but also one or more mid-level concepts are missing for the interpretation of the incoming resource. In case of missing explanations for mid-level concept instances, pattern P4 is always selected as prominent, to first enrich the ontology with missing mid-level concepts, thus enabling the interpretation of all mid-level concept instances. Then, in a subsequent cycle, the most suitable pattern will be chosen for handling the new situation appropriately.

3 Reasoning for Multimedia Interpretation

An ontology in a description logic framework is seen as a tuple consisting of a TBox and an ABox. In order to construct a high-level interpretation, the ABox part of the ontology is extended with some new assertions describing individuals and their relations. These descriptions are derived by media interpretation processes using the ontology (we assume the ontology axioms are denoted in a set Σ).

Interpretation processes are set up for different modalities, still images, videos, audio, and texts. In this section we discuss the interpretation process using an example interpretation for still images. The output is a symbolic description represented as an ABox. This ABox is the result of an abduction process (see [3] for a general introduction). In this process a solution for the following equation is computed: $\Sigma \cup \Delta \models \Gamma$. The solution Δ must satisfy certain side conditions.

The *abduction* inference service aims to construct a set of (minimal) ABox assertions Δ for a given set of assertions Γ such that Δ is consistent w.r.t. to the knowledge base $(\mathcal{T}, \mathcal{A})$ and satisfies [4]:

1. $T \cup A \cup \Delta \models \Gamma$
2. If $\Delta^{'}$ is an $ABox$ satisfying $T \cup A \cup \Delta^{'} \models \Gamma$, then $\Delta^{'} \models \Delta$ (Minimality)
3. $\Delta \not\models \Gamma$ (Relevance)

In Figure 3 an example from the athletics domain is presented. Assuming that it is possible to detect a horizontal bar bar_1, and a human $human_1$ by image analysis processes, the output of the analysis phase is represented as an ABox Γ. Assertions for the individuals and (some of) their relations detected by analysing Figure 3 are shown in Figure 4.

Fig. 3. Image displaying a high jump or a pole vault event.

$$\begin{aligned} bar_1 &: Horizontal\_Bar \\ human_1 &: Human \\ (bar_1, human_1) &: near \end{aligned}$$

Fig. 4. ABox Γ representing the result of the image analysis phase.

In order to continue the interpretation example, we assume that the ontology contains the axioms shown in Figure 5 (the ABox of the ontology is assumed to be empty). For some purposes, description logics are not expressive enough. Thus, some additional mechanism is required without jeopardizing decidability. In order to capture constraints among aggregate parts, we assume that the ontology is extended with DL-safe rules (rules that are applied to ABox individuals only). In Figure 6 a set of rules for the athletics example is specified. Note that the spatial constraints *touches* and *near* for the parts of a $Pole\_Vault$ event (or a $High\_Jump$ event) are not imposed by the TBox in Figure 5. Thus, rules are used to represent additional knowledge. Since spatial relations depend on the specific "subphases" of the events, corresponding clauses are included on the right-hand sides of the rules. For instance, a jumper as part of a $High\_Jump$ is near the bar if the image shows a $High\_Jump$ in the jump phase.

In the following we assume that rules such as those shown in Figure 6 are part of the TBox Σ. In order to provide a high-level interpretation, i.e. to provide a description of the image content in the form of high-level aggregates, we assume that spatial relations between certain objects detected by low-level analysis processes are not arbitrary.

In order to construct an interpretation for the image in Figure 3, an explanation is computed to answer why it is the case that a human is near a horizontal bar. Such explanations are considered the results of image interpretation processes. As mentioned above, the idea is to use the abduction inference service for deriving these kinds of (minimal) explanations (in the sense of interpretations). Minimal explanations can be extended appropriately in order to match expectations and task context.

We start with the computation of a minimal explanation in the athletics scenario. For this purpose, we slightly modify the abduction equation by taking into consideration that initially the ABox does not need to be empty. Thus, we divide Γ (see Figure 4)

$$
\begin{aligned}
Athlete &\equiv Human \sqcap \exists hasProfession.Sport \\
Jumper &\sqsubseteq Athlete \\
Pole &\sqsubseteq SportEquipment \\
Horizontal\_Bar &\sqsubseteq SportEquipment \\
Foam\_Mat &\sqsubseteq SportEquipment \\
Jumping\_Event &\sqsubseteq Event \sqcap \\
&\quad\ \exists_{\leq 1} hasParticipant.Jumper \\
Pole\_Vault &\sqsubseteq Jumping\_Event \sqcap \\
&\quad\ \exists hasPart.Pole \sqcap \\
&\quad\ \exists hasPart.Horizontal\_Bar \\
&\quad\ \exists hasPart.Foam\_Mat \\
High\_Jump &\sqsubseteq Jumping\_Event \sqcap \\
&\quad\ \exists hasPart.Horizontal\_Bar \\
&\quad\ \exists hasPart.Foam\_Mat \\
PV\_InStartPhase &\sqsubseteq \top \\
PV\_InEndStartPhase &\sqsubseteq \top \\
HJ\_InJumpPhase &\sqsubseteq \top
\end{aligned}
$$

Fig. 5. A tiny example TBox Σ for the athletics domain.

$$
\begin{aligned}
touches(Y, Z) \leftarrow\ & Pole\_Vault(X), \\
& PV\_InStartPhase(X), \\
& hasParticipant(X, Y), Jumper(Y), \\
& hasPart(X, Z), Pole(Z). \\
near(Y, Z) \leftarrow\ & Pole\_Vault(X), \\
& PV\_InEndStartPhase(X), \\
& hasPart(X, Y), Horizontal\_Bar(Y), \\
& hasParticipant(X, Z), Jumper(Z). \\
near(Y, Z) \leftarrow\ & High\_Jump(X), \\
& HJ\_InJumpPhase(X), \\
& hasPart(X, Y), Horizontal\_Bar(Y), \\
& hasParticipant(X, Z), Jumper(Z).
\end{aligned}
$$

Fig. 6. Additional restrictions for $Pole\_Vault$ and $High\_Jump$ in the form of rules.

into a part Γ_2 that the agent would like to have explained, and a part Γ_1 that the interpretation agent takes for granted. In our case Γ_2 is $\{(bar_1, human_1) : near\}$ and Γ_1 is $\{human_1 : Human, \quad bar_1 : Horizontal\_Bar\}$.

Coming back to the abduction problem specified above, we need solution(s) for the equation $\Sigma \cup \Delta \cup \Gamma_1 \models \Gamma_2$. In other words, given the background ontology Σ from Figures 5 and 6, the query

$$Q_1 := \{() \mid near(bar_1, human_1)\}$$

derived from Γ_2 should return $true$.

Obviously, this is not the case if Δ is empty. In order to see how an appropriate Δ could be derived, let us have a look at the rules in Figure 6. In particular, let us focus on the rules for $Pole\_Vault$ first. If we apply the rules to the query in a backward chaining way (i.e. from left to right) and unify corresponding terms we get variable bindings for Y and Z. The "unbound" variable X of the corresponding rules is instantiated with fresh individuals (e.g. pv_1). It is easy to see that two possible solutions for the abduction equation can be derived. For this example the output of the interpretation process are two interpretation ABoxes representing two possible interpretations of the same image (see Figures 7 and 8). As a result, this example illustrates the situation where evolution pattern P2 can be triggered.

$$
\begin{array}{rl}
human_1 : & Human \\
bar_1 : & Horizontal\_Bar \\
(bar_1, human_1) : & near \\
hj_1 : & High\_jump \\
hj_1 : & HJ\_InJumpPhase \\
human_1 : & Jumper \\
(hj_1, human_1) : & hasParticipant \\
(hj_1, bar_1) : & hasPart
\end{array}
$$

Fig. 7. ABox representing the first result of the abduction process.

$$
\begin{array}{rl}
human_1 : & Human \\
bar_1 : & Horizontal\_Bar \\
(bar_1, human_1) : & near \\
pv_1 : & Pole\_Vault \\
pv_1 : & PV\_InEndStartPhase \\
human_1 : & Jumper \\
(pv_1, human_1) : & hasParticipant \\
(pv_1, bar_1) : & hasPart
\end{array}
$$

Fig. 8. ABox representing the second result of the abduction process.

Note that due to the involvement of $human_1$ in the pole vault event in Figure 7, $human_1$ is now seen as an instance of $Jumper$, and, due to the TBox, also as an $Athlete$. Thus, information from high-level events also influences information that is available about the related parts. With queries for $Jumper$s the corresponding media objects would not have been found otherwise. Thus, recognizing high-level events is of utmost importance in information retrieval systems (and pure content-based retrieval does not help).

The example discussed here covers the interpretation of still images. It is necessary, however, to keep in mind that each media object might consist of multiple modalities, each of which will be the basis of modality-specific interpretation results (ABoxes). In order to provide for an integrated representation of the interpretation of media objects as a whole, these modality-specific interpretation results must be appropriately integrated. A cornerstone of this integration process will be to determine which modality-specific names refer to the same domain object. This will be discussed in later sections.

Continuing the example, it might be the case that for some images the ontology does not contain relevant axioms or rules. In this case, the interpretation result, i.e. the result of solving the abduction problem $\Sigma \cup \Delta \cup \Gamma_1 \models \Gamma_2$ will be degenerated because, due to missing axioms or rules in Σ, Δ must necessarily be equal to Γ_2 in order to solve the equation. As an example of such a situation we can discuss an interpretation of Figure 3 without the rules from Figure 6 and the GCIs for $Pole\_Vault$ and $High\_Jump$ in Figure 5. The degenerate interpretation result is shown (as Γ) in Figure 4. This result illustrates the situation where evolution pattern P4 can be triggered.

4 Ontology Population

Ontology population is the process of inserting concept instances and relation instances into an existing ontology. The proposed by BOEMIE approach for ontology population is built upon two axes: entity disambiguation and consistency maintenance. With entity disambiguation we refer to the process of identifying instances that refer to the same real object or event. If an ontology is populated with an instance without checking if the real object or event represented by the instance already exists in the ontology (as an instance that has populated the ontology at an earlier population step), then redundant information (in the best case) will be inserted into the ontology. A worst case scenario is the redundant instances to contain contradicting information, which may lead to an inconsistent ontology. At the same time maintaining the consistency of an ontology is crucial (mainly through the elimination of contradictory information), as an inconsistent ontology can not be used to reason with.

The ontology population activity can be decomposed into the following tasks:

- *Instance matching:* The first task of the population activity is the identification of similar instances contained in the ontology for each HLC instance (HLCi) in the set of explanations. Having a single HLCi as input, instance matching is expected to return a set of instances that populate the HLC and are similar to the incoming HLCi. Each returned matching instance is also expected to have a similarity figure, which measures the similarity of the two HLCis. The results of instance matching can be used to group instances that represent the same real object or event (as HLCis that are similar are assumed to represent the same real object or event) and possibly help in disambiguating multiple explanations in the case of evolution pattern P2 (during the ABox refinement task).
- *Instance grouping:* This task is responsible for grouping all the instances that represent the same real object or event, by exploiting the results of the instance matching task, where every incoming HLCi has been matched with a set of other instances that populate this HLC. Instance grouping is responsible to decide which of these matching instances clusters will be kept and grouped together to form a group that represents the same real object or event.
- *ABox refinement (evolution pattern P2 only):* In case of multiple explanations the most suitable explanation is selected by exploiting the results of the instance matching/grouping tasks. Assertions related to the rest of the explanations are removed from the ABox, thus leading to a refined version of it.
- *ABox validation:* This task performs consistency checking, to detect possible inconsistencies due to the additions that will be performed to the ontology. Two tasks must be validated: the addition of the incoming ABox (i.e., the one that originally triggered the population activity) into the ontology and the addition of a new instance of a grouping concept or the modification of an existing one. Both types of validation can be performed through standard reasoning (inference) services.
- *ABox assimilation:* The final task is responsible for performing the needed changes in the ontology (by creating all instances/relations in all ontological modules), in order to incorporate the information in the new ABox into the ontology.

Continuing the example presented in Section 3 where the ABox obtained by the semantic interpretation process contains two explanations (Figures 7 and 8), if this ABox is processed by the evolution pattern selector, then evolution pattern P2 will be selected, triggering a population operation. The first task of population (instance matching) will try to measure the similarity of each of the explanations found in the input ABox with instances already in the ontology, explained by the same concept. For example, in the case of the first explanation which is explained by an instance of the $High\_Jump$ concept (hj_2), instance matching will employ matching techniques in order to measure the similarity of the hj_2 instance with all instances that already populate the $High\_Jump$ concept. The same will happen for the pv_2 instance, and all other instances of HLCs that may be found in the input ABox (such as $human_2$). During instance grouping, all instances that were found similar during instance matching (i.e. they had a similarity above a certain threshold) will be checked whether they represent the same real object or event, and instances that do represent the same real object or event will be grouped by associating them with an instance that represents this real object or event (and no instances will be merged or eliminated from the ontology). For example, if there was enough information to identify that instance $human_2$ refers to the same person as another instance $human_x$ already in the ontology, then both $human_2$ and $human_x$ will be retained and associated with a new instance $human_{real}$ that represents the real person (while $human_2$ and $human_x$ are seen as "participations" of this person in specific event, location and time, which may even have different properties like age). Returning to our example, due to the very limiting information available, no grouping can be performed for any of the instances, as usually grouping requires information from more than one modality (i.e. the name of the athlete usually provided by the text modality is important to identify whether two instances refer to the same athlete).

Assuming that enough instances of both $High\_Jump$ and $Pole\_Vault$ concepts exist in the ontology, the ABox refinement task may be able to disambiguate multiple explanations, by selecting one instance as most "prominant". For example, the overal similarity of the hj_2 with other instances of the $High\_Jump$ concept may be better than the overal similarity of pv_2 with other instances of $Pole\_Vault$ (i.e. due to a missing pole instance). In such a case, a single explanation (i.e. hj_2) will be selected from the available explanations. In such a case, all instances not associated with hj_2, such as pv_2, will be discarded from the input ABox. Finally, during ABox validation the resulting ABox will be checked for consistency and if it is found to be consistent, it will be assimilated to the ontology.

5 Ontology Enrichment

Ontology enrichment is the activity of extending an ontology, through the addition of new concepts and relations. Ontology enrichment is performed every time the background knowledge is not sufficient to explain the extracted information from the processed multimedia documents. Thus, the ontology enrichment activity is expected to extend this background knowledge through the addition of new concepts/relations, in order to better explain extracted information in the future.

The ontology enrichment activity is triggered by either P3 or P4 evolution patterns. Evolution pattern P3 is selected when no explanation (i.e., an HLCi) has be found for a given ABox, and can lead to the insertion of a new HLC or a new relation into the ontology, or in the accumulation of the ABox in a "waiting" queue if available evidence cannot justify the addition of a new concept/relation. On the other hand, evolution pattern P4 is selected when the background knowledge is not sufficient to even assign MLCs to all of the extracted elements of a multimedia resource, thus inserting instances of the "unknown" MLC in the ABox. In this case, pattern P4 can result in the addition of a new MLC in the ontology. In fact, the detection of new MLCs is considered to have priority over the identification of new HLCs, because knowledge about a new MLC can lead to different semantic interpretation results about the same resources. As a result, when instances of the "unknown" MLC are found in ABoxes that contain no explanations, evolution pattern P4 is selected instead of P3.

Ontology enrichment is decomposed into the following tasks:

- *Concept learning:* The goal of this task is to propose new concepts (either HLCs or MLCs) and relations by exploiting similarities found through clustering, either in unexplained documents (evolution pattern P3) or in unknown objects recognised by the information extraction engine (evolution pattern P4). It can be decomposed into two main sub-tasks, clustering and concept formation.
  - *Clustering:* The main objective of the clustering task is to provide evidence that can support the creation of new concepts or relations.
  - *Concept formation:* This task is applicable only if a new HLC has been proposed by clustering. Exploiting the results of clustering, concept formation examines the clustered elements in order to extract common information (such as concepts/properties and relations) and use this common information to form a new concept, which is the result of this task.
- *Concept enhancement (evolution pattern P3 only):* This task is responsible for improving a concept identified by concept learning, through knowledge acquired from external sources, such as external domain ontologies or taxonomies.
- *Concept definition:* This task receives the new concept (either a new MLC or HLC) or relation as defined through the previous tasks, and shows the concept/relation definition to the ontology expert. The ontology expert must approve the new concept/relation in order to be assimilated into the ontology and additionally can revise the definition of the new concept/relation.
- *Concept validation:* This task performs consistency checking, by trying to detect possible inconsistencies due to the addition of the new concept relation to the ontology.
- *Concept assimilation:* The last task of ontology enrichment is responsible for performing the needed changes in the ontology in order to incorporate the newly formed concept/relation into the ontology.

As an example, we can assume an ABox where the semantic interpretation activity was unable to found an explanation, and in addition instances of the "unknown" MLC have been extracted by the information extraction toolkit. When such an ABox is processed by the evolution pattern selector, pattern P4 will be selected. In such a case, the ABox will be placed in a “waiting” stage. Once a significant number of instances of the “unknown” MLC have been assimilated, then clustering will be performed, during concept formation. If enough (modality specific) information is available that can lead to the formation of clusters, the concept definition task is responsible to present the results of the clustering to the ontology expert. The ontology expert must decide if the presented information is enough to justify the addition of a new MLC. In such a case, the expert must define the new concept(s), and associate all presented instances represented by the new concept with it.

In case of an ABox that has no explanation and also no instances of the “unknown” MLC, evolution pattern P3 will be selected. In such a case, the ABox will be also placed in a “waiting” stage, until enough ABoxes have been gathered. The ABox gathered are clustered in order to obtain clusters of ABoxes that are similar and all ABoxes in a cluster can possibly be explained by a single concept, which is not contained into the current version of the ontology. Once clusters have been identified, a new concept is formed for each found cluster by using all common information among all ABoxes of the cluster. Each newly formed concept will be further enhanced during concept enhancement, by trying to locate the formed concept in external knowledge sources through coordination and exploiting information from these knowledge sources, like concept/relation names and properties. Once new concepts/relations have been formed and possibly enhanced, they must be approved and reviewed by the ontology expert during the concept definition task: the expert must decide which of these proposed concepts/relations will be kept, what their definition will be and which ABoxes can be associated with them. Each concept/relation definition must be checked for consistency and assimilated into the ontology, if no inconcistencies have been found.

6 Original Contributions

The recent success of distributed and dynamic infrastructures for knowledge sharing has raised the need of semiautomatic/automatic ontology evolution strategies [5, 6]. An overview of some proposed approaches in this direction is presented in [7], even if limited concrete results have appeared in the literature. In most recent work, formal and logic-based approaches to ontology evolution are also being proposed. In [8], the authors provide a formal model for handling the semantics of change phase embedded in the evolution process of an OWL ontology. The proposed formalization allows to define and to preserve arbitrary consistency conditions (i.e., structural, logical, and user-defined).

A six-phase evolution methodology has been implemented within the KAON [9] infrastructure for business-oriented ontology management. The ontology evolution process starts with the capturing phase, that identifies the ontology modifications to apply either from the explicit business requirements or from the results of a change discovery activity. In the representation phase, the identified changes are described in a suitable format according to the specification language of the ontology to modify (e.g., OWL). The effects of the changes are evaluated in the semantics of change phase, where the ontology consistency check is also performed. Due to the fact that an ontology can reuse or extend other ontologies (e.g., through inclusion or mapping), the propagation phase ensures that any ontology change is propagated to the possible dependent artifacts in order to preserve the overall consistency. The subsequent implementation phase has the role to log all the performed changes in order to support the recovery facilities that in the final validation phase are provided to reverse an ontology change in case that an undesired effect occurs.

With respect to the state of the art literature on ontology evolution, original contributions of BOEMIE can be seen at two different levels, the whole methodology and the specific activities. The methodology as a whole proposes a new conceptualization of the problem of evolving multimedia ontologies, by presenting a pattern-driven evolution approach, where the most prominent evolution pattern for a specific evolution scenario is automatically identified on the basis of the results of the semantic interpretation activity against the background knowledge. Moreover, the methodology aims to minimize the human involvement by providing a set of learning, matching, and reasoning techniques that offer support in the various evolution activities, to allow the ontology expert to refine proposed working knowledge and/or to validate/choose among proposed alternative choices. Concerning novel contributions at the level of specific activities, the methodology for ontology population uses an innovative approach for the detection of instances which refer to the same real object or event, based on instance matching and non-standard clustering techniques. For ontology enrichment, clustering is used for detecting enough information to support the introduction of a new concept/relation and this supporting information is enhanced though information retrieved from external knowledge sources. Thus, the involvement of the ontology expert is reduced, as the expert is required to revise an already formed concept/relation than define this concept/relation from scratch. Matching techniques are used in combination with reasoning and clustering techniques in BOEMIE, thus leading to the development of a more flexible and comprehensive approach to concept/instance matching for evolution, by enforcing both structural matching and semantic matching.

The ontology evolution approach proposed by BOEMIE puts significant effort in maintaining the consistency of the ontology while trying on the same time to identify and eliminate redundant information. With respect to the state of the art in the field of knowledge representation and reasoning, the novel contribution of the evolution methodology regards the following issues: i) formalization of high-level multimedia interpretation as a logical decision problem and its implementation as a non-standard inference service, namely abduction; ii) extension of the knowledge representation formalism (SWRL) in order to support adequate knowledge representation for abduction tasks; iii) development of new optimization techniques for ABox consistency checking and query answering and its integration into a state-of-the art DL reasoner; iv) first approach for using non-standard description logic inference problems for formalizing the learning problem in BOEMIE (e.g., LCS, MSC, rewriting).

7 Concluding Remarks

In this paper, we have presented the methodology for multimedia ontology evolution developed in the context of the BOEMIE project. We have specified how the semantic interpretation of the information extracted from multimedia resources can be used to achieve the coordinated and consistent evolution of the ontology. We have described different evolution scenarios triggering population or enrichment patterns over the ontology. Current and future work within BOEMIE is devoted to: i) the investigation of instance matching techniques to support ontology population and to complement the role of reasoning techniques in the resource interpretation activity; ii) the investigation of clustering techniques to extract common information (such as concepts and relations) from ABoxes and use this common information to form new concepts; iii) the implementation of an ontology evolution toolkit providing an interactive environment where all the proposed techniques will be integrated into a coherent whole according to an open architecture.

References

1. Alani, H., Kim, S., Millard, D., Weal, M., Hall, W., Lewis, P., Shadbolt, N.: Automatic ontology-based knowledge extraction from web documents. IEEE Intelligent Systems **18** (2003) 14–21
2. Buitelaar, P., Cimiano, P., Racioppa, S., Siegel, M.: Ontology-based information extraction with soba. In: Proceedings of the International Conference on Language Resources and Evaluation. (2006) 2321–2324
3. Shanahan, M.: Perception as Abduction: Turning Sensor Data into Meaningful Representation. Cognitive Science **29** (2005) 103–134
4. Elsenbroich, C., Kutz, O., Sattler, U.: A case for abductive reasoning over ontologies. In: Proc. OWL: Experiences and Directions, Athens, Georgia, USA, November 10-11. (2006)
5. Haase, P., Sure, Y.: State-of-the-art on ontology evolution. SEKT informal deliverable 3.1.1.b, Institute AIFB, University of Karlsruhe (2004)
6. Klein, M., Noy, N.: A component-based framework for ontology evolution (2003)
7. Ding, Y., Foo, S.: Ontology research and development. part 2 - a review of ontology mapping and evolving. Journal of Information Science **28** (2002) 375–388
8. Haase, P., Stojanovic, L.: Consistent evolution of owl ontologies. In Gómez-Pérez, A., Euzenat, J., eds.: ESWC. Volume 3532 of Lecture Notes in Computer Science., Springer (2005) 182–197
9. Oberle, D., Volz, R., Motik, B., Staab, S.: An extensible ontology software environment. In Staab, S., Studer, R., eds.: Handbook on Ontologies. International Handbooks on Information Systems. Springer (2004) 311–333
10. Kim, S., Alani, H., Hall, W., Lewis, P., Millard, D., Shadbolt, N., Weal, M.: Artequakt: Generating tailored biographies from automatically annotated fragments from the web. In: Proceedings of Workshop on Semantic Authoring, Annotation & Knowledge Markup (SAAKM02), the 15th European Conference on Artificial Intelligence, (ECAI02), Lyon, France (2002) 1–6