

Bootstrapping Ontology Evolution with Multimedia Information Extraction

Georgios Paliouras, Constantine D. Spyropoulos, and George Tsatsaronis

Institute of Informatics and Telecommunications,
National Centre for Scientific Research “Demokritos”,
15310, Ag. Paraskevi, Attiki, Greece
{paliorug,costass,gbt}@iit.demokritos.gr

Abstract. This chapter summarises the approach and main achievements of the research project BOEMIE (Bootstrapping Ontology Evolution with Multimedia Information Extraction). BOEMIE introduced a new approach towards the automation of knowledge acquisition from multimedia content. In particular, it developed and demonstrated the notion of evolving multimedia ontologies, which is used for the extraction, fusion and interpretation of information from content of various media types (audio, video, images and text). BOEMIE adopted a synergistic approach that combines multimedia extraction and ontology evolution in a bootstrapping process. This process involves, on the one hand, the continuous extraction of semantic information from multimedia content in order to populate and enrich the ontologies and, on the other hand, the deployment of these ontologies to enhance the robustness of the extraction system. Thus, in addition to annotating multimedia content with semantics, the extracted knowledge is used to expand our understanding of the domain and extract even more useful knowledge. The methods and technologies developed in BOEMIE were tested in the domain of athletics, using large sets of annotated content and evaluation by domain experts. The evaluation has proved the value of the technology, which is applicable in a wide spectrum of domains that are based on multimedia content.

1 Motivation and Objectives of the BOEMIE Project

BOEMIE¹ aimed towards the automation of the knowledge acquisition process from multimedia content, which nowadays grows with increasing rates in both public and proprietary webs. Towards this end, it introduced the concept of *evolving multimedia ontologies*. The project was unique in that it linked multimedia extraction with ontology evolution, creating a synergy of great potential.

In recent years, significant advances have been made in the area of automated extraction of low-level features from audio and visual content. However, little progress has been achieved in the identification of high-level semantic features

¹ <http://www.boemie.org/>

or the effective combination of semantic features derived from various modalities. Driven by domain-specific multimedia ontologies, BOEMIE information extraction systems are able to identify high-level semantic features in image, video, audio and text and fuse these features for improved extraction. The ontologies are continuously populated and enriched using the extracted semantic content. This is a bootstrapping process, since the enriched ontologies in turn drive the multimedia information extraction system. Figure 1 provides a graphical illustration of this iterative bootstrapping process, that is implemented in the BOEMIE prototype. The main proposal of the project is illustrated by the continuous iteration that resides at the heart of the process. Information extraction is driven by semantic knowledge, while feeding at the same time the evolution of the ontologies.

Through the proposed synergistic approach, BOEMIE aimed at large-scale and precise knowledge acquisition from multimedia content. More specifically, the objectives of the project were:

Unifying representation for domain and multimedia knowledge. This multimedia semantic model follows modular knowledge engineering principles and captures the different types of knowledge involved in knowledge acquisition from multimedia. It realises the linking of domain-specific ontologies, which model salient subject matter entities, and multimedia ontologies, which capture structural and low-level content descriptions.

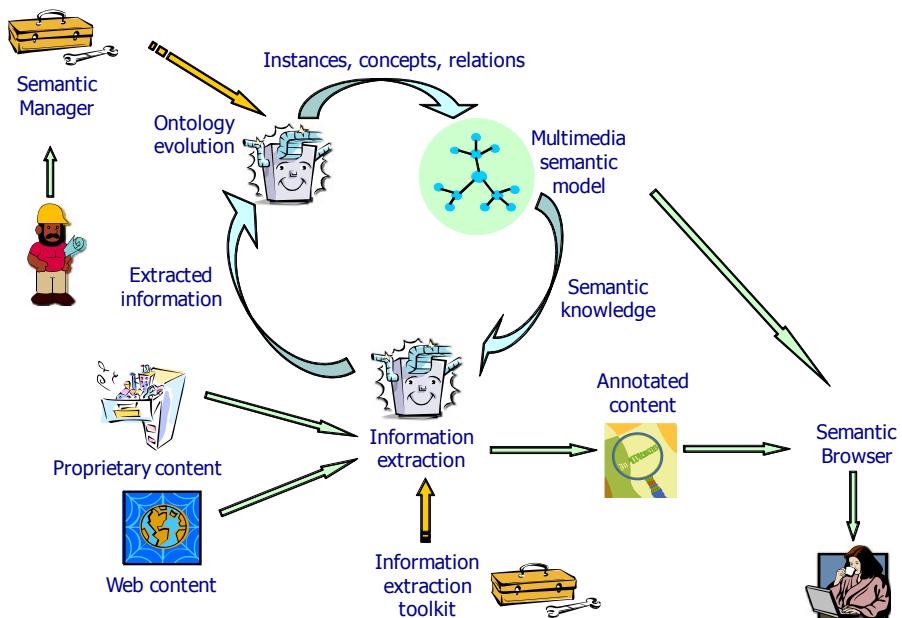


Fig. 1. The BOEMIE bootstrapping process

Methodology and toolkit for ontology evolution. The proposed methodology coordinates the various tools that use the extracted data to populate and enrich the ontologies. The toolkit provides tools to support ontology learning, ontology merging and alignment, semantic inference and ontology management.

Methodology and toolkit for information extraction. The methodology specifies how information from the multimedia semantic model can be used to achieve extraction from various media. Additionally, it fuses information extracted from multiple media to improve the extraction performance. The toolkit comprises tools to support extraction from image, audio, video and text, as well as information fusion.

The resulting technology has a wide range of applications in commerce, tourism, e-science, etc. One of the goals of the project was to evaluate the technology, through the development of an automated content collection and annotation service for athletics events in a number of major European cities. The extracted semantic information enriches a digital map, which provides an innovative and friendly way for the end user to access the multimedia content. Figure 2 illustrates this interaction of the end user with the system, which is provided by a specialised Web application, called the *BOEMIE semantic browser*. Points of interest that are associated with interesting multimedia content are highlighted on the map. The geo-referencing of the content is facilitated by the information extraction process of BOEMIE.

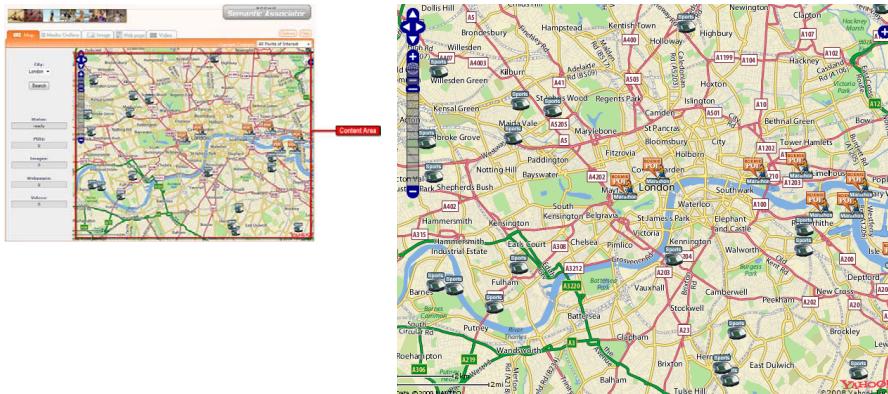


Fig. 2. The map-based interface to multimedia content in the semantic browser

The rest of this chapter is structured as follows. Section 2 presents briefly the main modules of the prototype system that was developed in BOEMIE. Section 3 compares BOEMIE to related projects that took place either before or in parallel with it. Finally, section 4 summarizes the main achievements of the project and proposes interesting paths for further research.

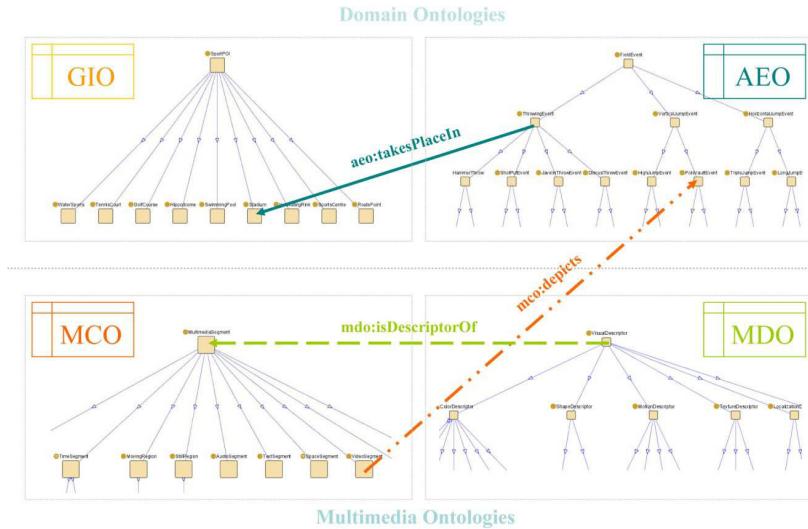


Fig. 3. The Multimedia Semantic Model. AEO (Athletics Event Ontology) models the scenario domain of interest, i.e. public athletics events. GIO (Geographic Information Ontology) models information relevant to geographic objects. MCO (Multimedia Content Ontology) models content structure descriptions, based on MPEG-7 MDS definitions. MDO (Multimedia Descriptor Ontology) models the MPEG-7 visual and audio descriptors.

2 The BOEMIE Prototype

More than 100 different modules and components have been produced in the course of the BOEMIE project, some of which have been made available publicly². Most of the components that were produced have been incorporated in the integrated prototype that was delivered and evaluated at the end of the project. The BOEMIE integrated prototype implements the bootstrapping process, as illustrated in Figure 1. This sketch shows also the main components of the prototype, which are described in the remaining of this section.

2.1 Multimedia Semantic Model

The BOEMIE Multimedia Semantic Model (MSM) [12,11] integrates ontologies that capture our knowledge about a particular domain, e.g. athletics, with ontologies that model knowledge about the structure and low-level descriptors pertaining to multimedia documents (Figure 3).

Besides addressing the interlinking of multimedia document segments with the corresponding domain entities, MSM further enhances the engineering of

² <http://www.boemie.org/software>

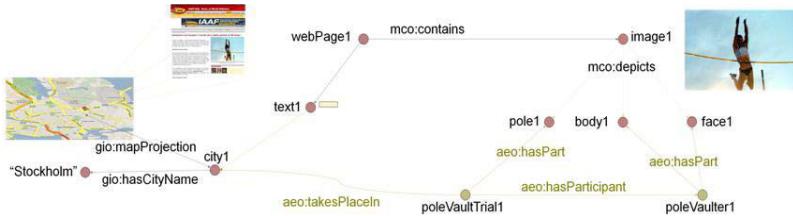


Fig. 4. Linking knowledge between ontologies in the semantic model

subject matter descriptions by distinguishing between *mid-level* (MLC) and *high-level* (HLC) domain concepts and properties, a feature unique to the BOEMIE project. Instances of MLCs represent information that is directly extracted from the multimedia content, using the various analysis tools, e.g. the name of an athlete or her body in a picture. On the other hand, instances of HLCs are generated through reasoning-based interpretation of the multimedia content, using the domain ontology. Such engineering allows incorporating the analysis perspective into the domain conceptualisation, which in turn supports effective logic-grounded interpretation. The developed ontologies allow the utilisation of precise formal semantics throughout the chain of tasks involved in the acquisition and deployment of multimedia content semantics.

In MSM, four OWL DL ontologies are linked in a way that supports the purposes of BOEMIE for semantics extraction, interpretation, evolution, as well as retrieval and representation of the acquired semantics. Figure 4 presents a simple example of this interlinking between the ontologies. This is also the main novelty of the BOEMIE Multimedia Semantic Model.

2.2 Recursive Media Decomposition and Fusion

The information extraction toolkit of BOEMIE integrates a number of tools for content analysis and interpretation, using a recursive media decomposition and fusion framework. In the course of the project, innovative methods for the analysis of single-modality content were developed, going in most cases beyond the state-of-the-art. As Figure 5 illustrates, these methods cover most of the currently available types of media. Most importantly, they support the bootstrapping process through an evolving cycle of analysis of new content, learning of improved analysis models and discovering interesting objects and entities to add to the domain knowledge.

The coordination of the evolving extraction process is achieved by a new method that was developed in BOEMIE and is called Recursive Media Decomposition and Fusion (RMDF) [21]. The method decomposes a multimedia document into its constituent parts, including embedded text in images and speech. It then relies on single-modality modules, the results of which are fused together in a common graph that complies with the domain ontology. In a final step, graph techniques are used to provide a consistent overall analysis of the multimedia

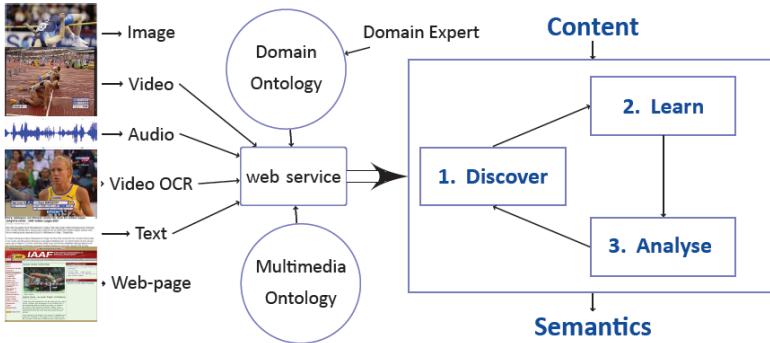


Fig. 5. The Information Extraction Toolkit

document. For instance, a Web page may be decomposed into several parts, one of which containing a video, which may in turn contain a static overlaid image, that embeds text, which refers to a person. This example shows the importance of recursive decomposition and corresponding fusion of results that come from video and text analysis.

Regarding the single-modality modules, BOEMIE has developed innovative methods to:

- detect and discover objects of various shapes and sizes in images [22],
- track moving objects in video and classify movement phases,
- identify and discover entities in text and relations amongst them [14,20],
- detect overlay and scene text in video and perform optical character recognition on it [23,1],
- recognise and discover audio events and interesting keywords in audio [4,5].

Figure 6 provides examples of such results. Most importantly, however, through interpretation and fusion, the RMDF is able to improve significantly the precision of multimedia analysis, be it in Web pages containing HTML text and images or video footage with audio commentary and overlay text. In addition to the novel decomposition and fusion approach, the single-modality tools support customization to any domain, by allowing the discovery of new semantics in content and learning to identify known objects and entities. Furthermore, the extraction toolkit is easily distributable and scalable, by dynamically integrating per media analysis techniques in an unrestricted number of servers, communicating through a computer network.

2.3 Abductive Multimedia Interpretation

The interpretation of multimedia content in BOEMIE goes well beyond the usual extraction of semantics from individual media. Domain knowledge, in the form of ontologies, is being exploited by a reasoning-based interpretation service that



(a) Object detection in images.



(b) Phase detection in video.



(c) Text recognition in video.



(d) Entity and relation extraction in Web pages.

Fig. 6. Sample results of single-media analysis tools

operates in two levels: single-media interpretation and fusion. The interlinking of domain and multimedia ontologies in the semantic model (Fig. 4) support this process. Figure 7 illustrates the multi-level analysis and interpretation process. Both the single-media and the fusion services are supported by the same reasoning apparatus.

Reasoning for multimedia interpretation is based on the RacerPro reasoning engine³, which has been extended with many novel methods for the purposes of BOEMIE [3,17,16]. One of the main extensions is the use of abduction to generate interpretation hypotheses for what has been “observed” by the extraction tools. The new abductive query answering service of RacerPro is able, during query evaluation, to recognize non-entailed query atoms and hypothesize them. Since there might be more than one hypothesis (i.e. explanations), a set of

³ <http://www.racer-systems.com/products/racerpro/>

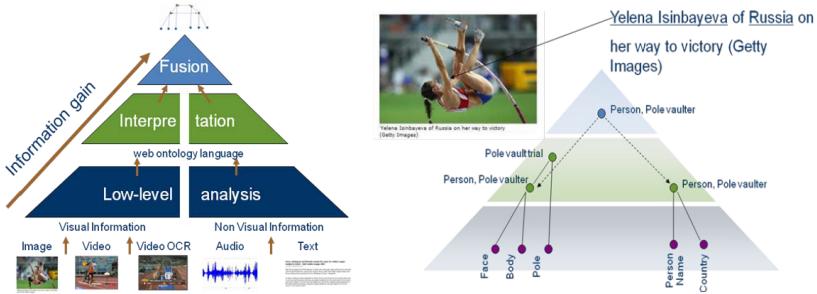


Fig. 7. Multimedia analysis and interpretation process

scoring functions has been designed and implemented in order to prefer certain hypotheses over others. Given the complexity of the interpretation hypothesis (a.k.a. explanation) space, important optimizations have been developed in the reasoner, in order to cut down on the number of consistent and useful interpretations that are produced by the system.

The novel abductive multimedia interpretation machinery of BOEMIE combines Description Logics, as a representation formalism for ontologies, with DL-safe rules that guide the search for interpretations. In the context of BOEMIE, methods to learn these rules have also been developed.

2.4 Pattern-Based Ontology Evolution

The ontology evolution toolkit (OET) implements a pattern-based approach to the population and enrichment of the ontology, which is unique to BOEMIE [10]. In particular, two different cases have been identified for the population process, one in which a single interpretation is produced for a document and one in which more than one candidate interpretation is provided. Furthermore, two cases are defined for the enrichment process, one in which a high-level concept (HLC) and one in which a mid-level concept (MLC) is added. Each of those cases requires different handling in terms of the interaction with the domain expert and the modules that are employed for the semi-automated generation of new knowledge, e.g. concept enhancement, generation of relations and interpretation rules, etc. Figure 8 provides a high-level overview of these four cases (patterns P1 to P4) and the modules that are involved.

The first two cases (P1 and P2), which are responsible for the population of the ontology with new instances, are primarily based on instance matching and grouping methods[7,9]. Novel methods have been developed for this purpose, in order to take advantage of the rich semantics of the BOEMIE semantic model and scale efficiently to large document sets. These methods have been incorporated in the HMatch ontology matching software [6], which is publicly available⁴.

A number of innovations have been made also in the area of ontology enrichment (patterns P3 and P4) [18,15]. The discovery of new concepts and properties

⁴ [tt <http://islab.dico.unimi.it/hmatch/>](http://islab.dico.unimi.it/hmatch/)

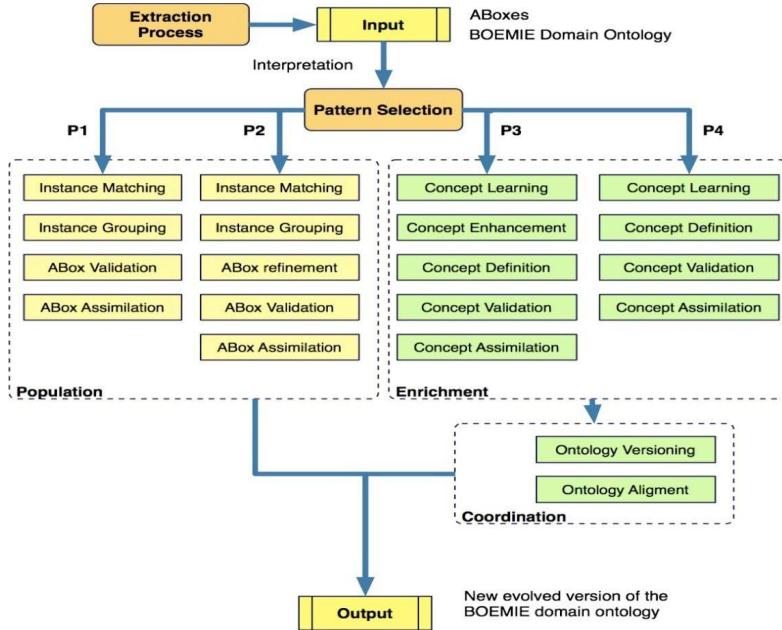


Fig. 8. Pattern-based ontology evolution

is based on a new methodology that incorporates a set of ontology modification operators. Logical and statistical criteria are introduced for the choice of the most appropriate modifications to the ontology, given the observed data. Further to this data-driven enrichment, a concept enhancement method has been developed, matching new constructs to knowledge in external resources, e.g. on the Web.

2.5 Interface Components

In addition to the core processing components, the BOEMIE prototype includes a number of interface components that facilitate the interaction of the users with the prototype, as well as the interaction among the components. Three of these components, which introduce a number of novel features are the *Semantic Browser*, the *Semantic Manager* and the *Bootstrapping Controller*.

The *BOEMIE Semantic Browser (BSB)* provides an innovative interaction experience with multimedia content. It does so by supporting three modes of interaction with the multimedia content:

- Interactive maps for multimedia retrieval.
- Interactive content of media objects.
- Dynamic suggestion of related information.

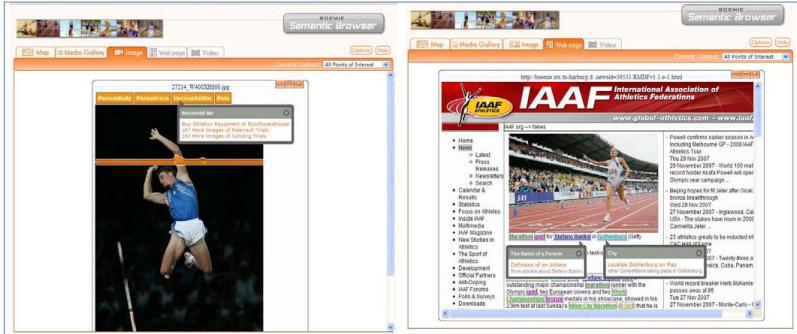


Fig. 9. Suggesting information related to active media objects

A screenshot of the map interface of BSB is shown in Fig. 2. Based on the information extraction technologies of BOEMIE, BSB can associate multimedia content with geopolitical areas and specific Points of Interest (PoIs) on digital maps. In this manner it provides direct access to the multimedia, through what we call “BOEMIE PoIs”. Furthermore, BSB uses the semantic annotations generated automatically by BOEMIE to make media objects interactive. More specifically, it automatically highlights relevant content of a specific domain on top of text or images to prepare the interface for further interaction possibilities. Finally, through interpretation, BOEMIE is able to generate deeper semantic information, e.g. the type of sport that an image depicts. Using this implicit knowledge, BSB provides context-sensitive advertisement and suggests related information. This is realized by the idea of context menus, illustrated in Fig. 9.

The BOEMIE Semantic Manager (BSM) [8] is unique in its simplification of a complex and demanding process, i.e., that of adding semantics to multimedia content and maintaining the associated domain knowledge. As an interface to the OET, BSM provides three primary functionalities:

- Population of the ontology with semantically annotated multimedia content.
- Enrichment of the ontology with new knowledge that has been learned from data.
- User-friendly interactive enhancement of new knowledge by the domain expert.

BSM provides interactive selection/approval/rejection of the multimedia content interpretations automatically produced by the BOEMIE system, as well as (similarity-based) document browsing facilities. In order to make the process accessible to the non-skilled in knowledge engineering, it creates a natural language description of the underlying logic representation of ontology instances. Additionally, BSM provides terminological and structural suggestions to support the domain expert in performing ontology enrichment. Suggestions are dynamically extracted from knowledge chunks similar to a given concept proposal by

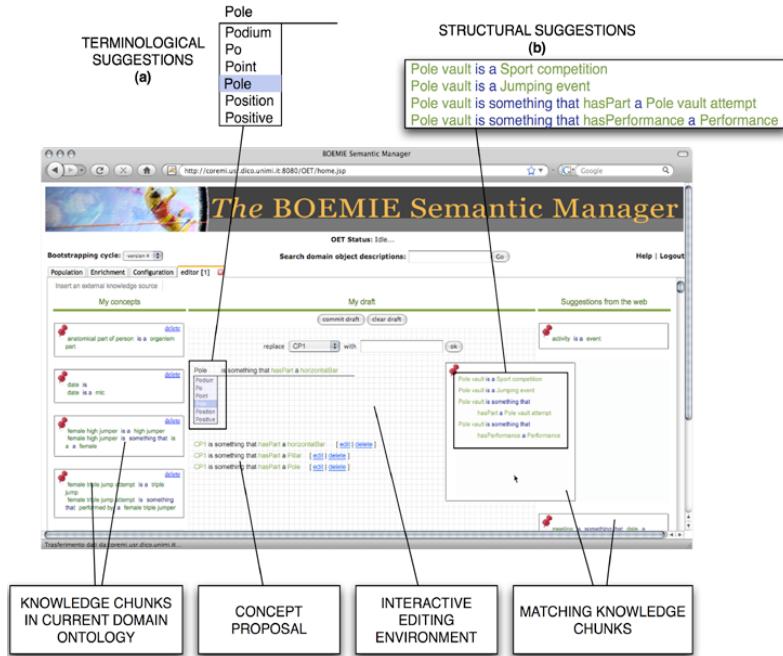


Fig. 10. Concept definition and enhancement, using the semantic manager

exploiting ontology matching techniques. A repository of knowledge chunks is created and maintained through a knowledge harvesting process that periodically searches knowledge of interest from other ontologies, web directories, and, in general, external knowledge repositories. Finally, BSM incorporates a simple ontology editor [13], which uses natural language patterns and autocomplete techniques to facilitate the incorporation of new knowledge to the domain ontology. Figure 10 illustrates this editor.

The Bootstrapping Controller (BSC) is the main application logic component that implements the iterative extraction and evolution process. Using the BSC, the content owner can add content to the Multimedia Repository (MMR) and then send it for processing through predefined workflows. The content is added to the repository, either by uploading specific files or by crawling the Web. Typically, a new document will be sent to RMDF for extraction and the results of its interpretation will be populated into the ontology. When sufficient evidence is accumulated, the OET will generate proposals for changes to the ontology. The domain expert will use these recommendations to change the ontology and the content will be sent again for processing by the RMDF. In some cases, new mid-level concepts (MLCs) will be generated based on the analysis of the multimedia content so far. In these cases, in addition to the extension of the ontologies, the BSC will send sufficient training data to the RMDF, asking for the re-training of the analysis modules.

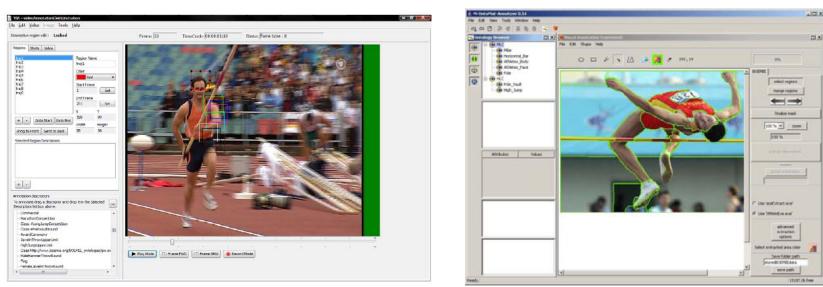


Fig. 11. Video and image annotation with the VIA tool

2.6 Manual Annotation Tools

The information extraction methods developed in BOEMIE are trainable and therefore require training material, in order to learn to identify interesting entities, objects and relations among them in multimedia content. The BOEMIE bootstrapping process generates semi-automatically such training data. However, for the purposes of training and evaluating the initial extractors, we generated significant quantities of training data for all types of media: image, video, audio, text. For these data we used interactive tools for manual annotation. Most of these tools were also developed in BOEMIE and improve significantly the state of the art in the field.

The VIA tool⁵ can be used for high-level and low-level video and image annotation. In both cases, annotation is aligned with concepts of the domain ontologies. In the case of image annotation, either image regions and complete images are linked with concepts (high-level annotation) or visual descriptors are extracted per annotated region and associated with the corresponding concept (low-level annotation). To reduce the manual annotation burden, VIA supports the automatic segmentation of a still image into regions and region-merging. Regarding video annotation, VIA supports input in MPEG1/2 video format and frame accurate video playback and navigation. Video annotation can take place either in a frame-by-frame style or as live annotation during playback. Figure 11 illustrates the use of VIA.

The text and HTML annotation tool BTAT⁶ [19] has been developed over the Ellogon open-source text engineering platform⁷. It supports the annotation of named entities, the mid-level concepts (MLCs), as well as relations between those named entities. The relations are grouped in tables of specific types. Tables correspond to high-level concepts (HLCs). Furthermore, the tool enables the annotation of relations between HLC instances by creating links between tables in an effective and easy way. One of the innovations in BTAT is its dual manual and automated annotation functionality. Manual annotation is facilitated by a smart

⁵ <http://mklab.iti.gr/project/via>

⁶ <http://www.boemie.org/btat>

⁷ <http://www.ellogon.org/>

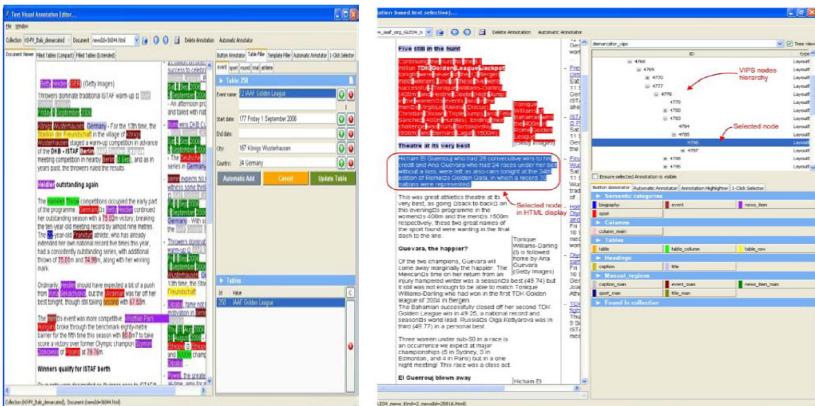


Fig. 12. Text and HTML annotation with the BTAT tool

text-marking system, where the user selects with a mouse click words, instead of single characters. Automatic annotation works by matching user-defined regular expressions. Figure 12 illustrates the use of BTAT.

3 Related Projects

BOEMIE is part of a larger research effort to provide easier access to the ever increasing quantities of multimedia content, particularly on the Web. As such it was preceded and followed by a number of related efforts that influenced or were influenced by it. An extensive list of related projects provided on the BOEMIE Web site⁸. In this section, we refer to a selection of these projects that we consider most relevant.

BOEMIE brought together experts from a number of research fields, who have worked on the analysis of textual, visual and audio content in the past. Even before the beginning of BOEMIE, some of these efforts aimed at the extraction of high-level information, i.e. semantics, from such content. One of the early efforts in this direction was the SCHEMA project⁹, which developed technology for content-based analysis and retrieval of multimedia, using ontologies. SCHEMA was followed by the AceMedia project¹⁰, which went further into defining ontologies for multimedia and developing an initial version of a manual annotation tool for images. In parallel, the CROSSMARC project¹¹ developed ontology-based information extraction technology for cross-lingual text documents, including initial attempts to introduce ontology evolution methods into the process.

⁸ http://www.boemie.org/sites/default/files/BOEMIE_related_projects_with_contacts.pdf

⁹ <http://www.iti.gr/SCHEMA/>

¹⁰ <http://cordis.europa.eu/ist/kct/acemediaSynopsis.htm>

¹¹ <http://labs-repos.iit.demokritos.gr/skel/crossmarc/>

Apart from the projects in which BOEMIE partners participated actively, there was a range of other projects on which BOEMIE was based. European networks of excellence, focusing primarily on infrastructure work, were among the major sources of knowledge for BOEMIE. In particular, the OntoWeb project¹² and its follow-up KnowledgeWeb¹³ provided interesting survey reports and resources on semantic Web research, which were valuable in the early stages of BOEMIE. Furthermore, the MUSCLE project¹⁴ was a useful source of information about multimedia semantics extraction technology. Apart from the networks of excellence, two European projects that run in parallel to BOEMIE are worth-mentioning: the SEKT¹⁵ and the ALVIS¹⁶ projects. SEKT emphasized on semantic Web ontologies, including ontology matching and enrichment, while ALVIS focussed on semantic search. Beyond European research, there were a number of other projects, particularly in the US, that have provided useful input to BOEMIE. Among these, the projects Marvel¹⁷ and IMKA are particularly worth-mentioning. MARVEL was an IBM research project that aimed at large-scale classification of visual content, while IMKA [2] was a project by the University of Columbia, emphasizing knowledge representation for multimedia.

A number of other interesting projects either run in parallel or followed BOEMIE. Members of the BOEMIE consortium contributed to some them. The Mesh project¹⁸ developed ontology-based extraction technology for multimedia news content, while the LIVE project¹⁹ focused on real-time video search and editing for news broadcasting. In the same industry, CASAM²⁰ develops technology for computer-assisted annotation of video content by news editors. On the other hand, the WeKnowIT project²¹ advocates the importance of social generation of multimedia content, which is centered around important events. Finally, following the series of networks of excellence, the K-Space project²² produces useful infrastructure for ontology-based analysis of multimedia content.

Apart from the projects in which BOEMIE partners participate, a range of other projects, related to BOEMIE have started. Among these, the projects X-media²³, Vidi-Video²⁴ and Vitalas²⁵ aimed to move into large-scale semantic indexing and retrieval of multimedia content, in particular image and video.

¹² <http://www.ontoweb.org/>

¹³ <http://cordis.europa.eu/ist/kct/knowledgewebSynopsis.htm>

¹⁴ <http://cordis.europa.eu/ist/kct/muscleSynopsis.htm>

¹⁵ <http://cordis.europa.eu/ist/kct/sektSynopsis.htm>

¹⁶ <http://cordis.europa.eu/ist/kct/alvisSynopsis.htm>

¹⁷ <http://www.research.ibm.com/marvel/>

¹⁸ <http://www.mesh-ip.eu/?Page=Project>

¹⁹ <http://www.ist-live.org/>

²⁰ <http://www.casam-project.eu/>

²¹ <http://www.weknowit.eu/>

²² <http://cordis.europa.eu/ist/kct/kspaceSynopsis.htm>

²³ <http://cordis.europa.eu/ist/kct/x-mediaSynopsis.htm>

²⁴ <http://cordis.europa.eu/ist/kct/vidivideoSynopsis.htm>

²⁵ <http://vitalas.ercim.org/>

On the other hand, the Bootstrep project²⁶ studies new methods for the continuous evolution of ontologies and LarKC²⁷ aims to provide reasoning technology that could be used with large knowledge bases.

4 Summary and Open Issues

The BOEMIE project has brought tightly together two complementary technologies, namely information extraction from multimedia and ontology evolution. It has done so, by introducing a bootstrapping framework, within which each of the two technologies iteratively feeds the other. In this framework, multimedia content of all known types is analysed and interpreted, using multimedia ontologies. The results of multimedia interpretation are used to populate the ontologies and initiate the enrichment of domain knowledge. The improved ontologies can then be used to re-analyse the multimedia content and extract additional information that is translated into new knowledge.

We consider the methods and technology developed in BOEMIE an important contribution to the field of semantic analysis of multimedia. More recent work in this field has emphasized the scalability of the process, aiming to make the corresponding technologies applicable to as large sets of content as the Web itself. The move from the old Web to its social counterpart has intensified the need for such technology, due to the unprecedented volume of media generation by non-expert users. It is thus to be expected that adaptive solutions, such as those proposed by BOEMIE, will play an increasingly important role in the future of multimedia analysis and search.

Acknowledgements

BOEMIE was a collaborative project that was partially funded by the European Commission, under contract number FP6-IST-027538. The work presented here is the result of the common effort of more than 30 people from the following six research teams: Institute of Informatics and Telecommunications (National Centre for Scientific Research “Demokritos”, Greece), Institute for Intelligent Analysis and Information Systems (Fraunhofer Gesellschaft, Germany), Institute of Software Systems (Hamburg University of Technology, Germany), Department of Informatics and Communication (Università degli Studi di Milano, Italy), Informatics and Telematics Institute (Centre for Research and Technology - Hellas, Greece), Tele Atlas N.V. (The Netherlands).

References

1. Anthimopoulos, M., Gatos, B., Pratikakis, I.: Multiresolution text detection in video frames. In: Proceedings of the 2nd International Conference on Computer Vision Theory and Applications (VISAPP), Barcelona, Spain, pp. 161–166 (March 2007)

²⁶ <http://cordis.europa.eu/ist/kct/bootstrap-synopsis.htm>

²⁷ <http://www.larkc.eu/>

2. Benitez, A.B., Chang, S.-F., Smith, J.R.: Imka: A multimedia organization system combining perceptual and semantic knowledge. In: Proceedings of the Ninth ACM International Conference on Multimedia (MULTIMEDIA), pp. 630–631 (2001)
3. Berger, T., Kaplunova, A., Kaya, A., Möller, R.: Towards a scalable and efficient middleware for instance retrieval inference services. In: Proceedings of the 3rd International Workshop on OWL: Experiences and Directions (OWLED), co-located with the 4th European Semantic Web Conference (ESWC 2007), Innsbruck, Austria (June 2007)
4. Biatov, K.: Two level discriminative training for audio event recognition in sport broadcasts. In: Proceedings of the 12th International Conference on Speech and Computer (SPECOM), Moscow, Russia (October 2007)
5. Biatov, K., Hesseler, W., Kohler, J.: Audio data retrieval and recognition using model selection criterion. In: Proceedings of the 2nd International Conference on Signal Processing and Communication Systems, Gold Coast, Australia (December 2008)
6. Castano, S., Ferrara, A., Lorusso, D., Montanelli, S.: The hmatch 2.0 suite for ontology matchmaking. In: Proceedings of the 4th Italian Workshop on Semantic Web Applications and Perspectives (SWAP), Bari, Italy (December 2007)
7. Castano, S., Ferrara, A., Lorusso, D., Montanelli, S.: On the ontology instance matching problem. In: Proceedings of the 19th International Conference on Database and Expert Systems Applications (DEXA), pp. 180–184 (2008)
8. Castano, S., Ferrara, A., Montanelli, S.: Evolving multimedia ontologies: the bsm tool environment. In: Proceedings of 17th Italian Symposium on Advanced Database Systems (SEBD), Camogli, Italy (June 2009)
9. Castano, S., Ferrara, A., Montanelli, S., Lorusso, D.: Instance matching for ontology population. In: Proceedings of the Sixteenth Italian Symposium on Advanced Database Systems (SEBD), Mondello, Italy (June 2008)
10. Castano, S., Peraldi, S.E., Ferrara, A., Karkaletsis, V., Kaya, A.: Multimedia interpretation for dynamic ontology evolution. Journal of Logic and Computation 19(5), 859–897 (2009)
11. Dasiopoulou, S., Tzouvaras, V., Kompatsiaris, I., Strintzis, M.G.: Capturing mpeg-7 semantics. In: International Conference on Metadata and Semantics Research (MTSR), Corfu, Greece, pp. 113–122 (2007)
12. Dasiopoulou, S., Tzouvaras, V., Kompatsiaris, I., Strintzis, M.G.: Enquiring mpeg-7 based ontologies. Multimedia Tools and Applications 46(2), 331–370 (2010)
13. Ferrara, A., Montanelli, S., Varese, S.C.G.: Ontology knowledge authoring by natural language empowerment. In: IEEE Proceedings of the 1st DEXA Int. Workshop on Modelling and Visualization of XML and Semantic Web Data (MoViX), Linz, Austria (September 2009)
14. Fragkou, P., Petasis, G., Theodorakos, A., Karkaletsis, V., Spyropoulos, C.D.: Boemie ontology-based text annotation tool. In: Proceedings of the 6th International Conference on Language Resources and Evaluation (LREC), Marrakech, Morocco (May 2008)
15. Peraldi, S.E., Kaya, A.: Ontology and rule design patterns for multimedia interpretation. In: Yamaguchi, T. (ed.) PAKM 2008. LNCS (LNAI), vol. 5345. Springer, Heidelberg (2008)
16. Peraldi, S.E., Kaya, A., Melzer, S., Möller, R.: On ontology based abduction for text interpretation. In: Gelbukh, A. (ed.) CICLing 2008. LNCS, vol. 4919, pp. 194–205. Springer, Heidelberg (2008)

17. Peraldi, S.E., Kaya, A., Melzer, S., Möller, R., Wessel, M.: Towards a media interpretation framework for the semantic web. In: Proceedings of the IEEE/WIC/ACM International Conference on Web Intelligence (WI), Silicon Valley, USA, pp. 374–380 (November 2007)
18. Petasis, G., Karkaletsis, V., Krithara, A., Palouras, G., Spyropoulos, C.D.: Semi-automated ontology learning: the boemie approach. In: Proceedings of the First ESWC Workshop on Inductive Reasoning and Machine Learning on the Semantic Web, Heraklion, Greece (June 2009)
19. Petasis, G., Fragkou, P., Theodorakos, A., Karkaletsis, V., Spyropoulos, C.D.: Segmenting html pages using visual and semantic information. In: Proceedings of the 4th Web as Corpus Workshop: Can we do better than Google?, at the 6th International Conference on Language Resources and Evaluation (LREC), Marrakech, Morocco (May 2008)
20. Petasis, G., Karkaletsis, V., Palouras, G., Spyropoulos, C.D.: Learning context-free grammars to extract relations from text. In: Proceedings of the 18th European Conference on Artificial Intelligence (ECAI), Patras, Greece, pp. 303–307 (July 2008)
21. Petridis, S., Perantonis, S.J.: Semantics extraction from multimedia data: an ontology-based machine learning approach. In: Cutsuridis, A.H.V., Taylor, J.G. (eds.) Perception-action Cycle: Models, Architectures and Hardware (2011)
22. Tsapatsoulis, N., Petridis, S., Perantonis, S.J.: On the use of spatial relations between objects for image classification. In: Proceedings of the 4th IFIP Conference on Artificial Intelligence Applications and Innovations (AIAI), Athens, Greece (September 2007)
23. Vamvakas, G., Gatos, B., Pratikakis, I., Stamatopoulos, N., Roniotis, A., Perantonis, S.J.: Hybrid off-line ocr for isolated handwritten greek characters. In: Proceedings of the 4th IASTED International Conference on Signal Processing, Pattern Recognition, and Applications (SPPRA), Innsbruck, Austria (February 2007)