AGENTS AND ONTOLOGIES FOR UNDERSTANDING AND PRESERVING THE ROCK ART OF MOUNT BEGO

L. Papaleo
ICT Department, Province of Genova, Italy

G. Quercini
Institute for Advanced Computer Studies, University of Maryland, USA

V. Mascardi†, M. Ancona†, A. Traverso‡
† CS Department, ‡ DARFICLET, University of Genova, Italy

H. De Lumley
Laboratoire Départemental de Préhistoire du Lazaret, ADEVREPAM, Nice, France

Keywords: MAS for rock-art preservation, ontology extraction, sketch interpretation, visual database for rock-art.

Abstract: This paper describes the joint effort of computer scientists, archaeologists, and historians for designing a multi-agent system that exploits ontologies for the semantic description of the Mount Bego petroglyphs, thus moving a step forward their preservation. Most components of the MAS have already been developed and tested, and their integration is under way.

1 INTRODUCTION

For decades the area around Mount Bego has been deemed as a sort of a bewitched place, rocks being carved with “thousand devils” (Pierre de Montfort, XV century). On the other hand, archaeologists and historians look at this place as an incredibly valuable source of knowledge, due to the up to 40,000 figurative petroglyphs and 60,000 non-figurative petroglyphs scattered over a large area at an altitude of 2,000 to 2,700 meters.

The historical relevance of the Mount Bego petroglyphs is unquestionable, as they date back to the early Bronze Age, when humans left no written evidences and the only witnesses of their existence are their tools and, indeed, their “drawings”.

While learning from written sources is relatively easy, even if this may actually depend on the source, images generally lend themselves to a number of different and often conflicting interpretations, which is the case of the Mount Bego carved rocks. Hence, the coherence of any new interpretation of petroglyphs must be checked against multiple sources.

Another major issue is that Mount Bego rocks are not protected in a safe place such as a museum and thus they are constantly exposed to rough weather as well as vandalism of careless or malicious visitors.

The explorer who first realized the importance of Mount Bego carvings from an historical point of view and started a systematic activity for preserving them, was Clarence Bicknell. In 1897, he sketched 450 drawings on small sheets of paper. Between 1898 and 1910 he realized up to 13,000 drawings and reliefs, part of which were then published in (Bicknell, 1913). Bicknell’s collection is completed by inedited drawings owned by the University of Genoa. These amount to 16,000 drawings and reliefs on different materials. At the foot of almost all sheets Bicknell wrote personal notes on the depicted subject, the location of the petroglyph, a name assigned to the rock with the petroglyph and the date of the relief. The legacy owned by the University of Genoa also includes nine notebooks, filled with notes in Victorian English, which can be subdivided into five excavation diaries and four note books. which cover a timespan of 10 years (1902-1912).

Many years after Bicknell’s campaigns, several teams led by Henry de Lumley have been surveying and mapping this important archaeological area starting from 1967.

To the purpose of preserving Mount Bego rock art, we are moving along two directions: first, we will integrate the wonderful collection of Bicknell into an existing database of data relative to Mount Bego currently hosted by the Laboratoire Départemental de Préhistoire du Lazaret (Adevrepam), based in Nice, France; then we have implemented, and we are currently integrating into our MAS, most of the methods and interfaces that will allow end users to add other data that might become available in the future:

data will be recognized and semi-automatically categorized, according to specific ontology-driven criteria.

The reminder of this paper is organized as follows. An overview of our approach is discussed in Section 2; details on the design and partial implementation of a multi-agent system supporting sketch interpretation algorithms are given in Section 3; semi-automatic extraction of ontologies from texts is discussed in Section 4. In Section 5 we describe how our framework will be used to integrate, analyze and interpret data on Mount Bego. Related work and concluding remarks are finally presented in Section 6.

2 Overview of our Approach

The cultural relevance of the Mount Bego area calls for immediate action for preservation. For this reason the domain experts involved in this paper (H. de Lumley himself and A. Traverso, an historian working at DARFICLET, the Department of Archaeology, Classical Philology and their Traditions in the Christian, Medieval and Humanistic Ages, University of Genoa) welcomed the computer scientists' proposal to complement the techniques they usually adopt to preserve the area with digital preservation and restoration techniques.

Our long-term goal is the creation of a repository which may eventually become a reference at European level as a repository for Bronze Age petroglyphs. To this aim the repository needs to come with tools that ease the integration of new data and allow analysis, comparison, interpretation and search on them. Although here we only focus on Mount Bego, our approach is general enough to handle scenarios other than Mount Bego. Just to make an example, many studies[1] demonstrate that there are strong similarities between Mount Bego and Valcamonica's petroglyphs. Hence, our approach should easily apply to preserve and maintain in digital form Valcamonica's petroglyphs, in the same way as it applies to Mount Bego's ones, coherently with our long-term objective.

In the current setting we have access to a database managed by Adevrepam, whose data are the best candidates for starting to fill the European Bronze Age petroglyphs repository. It is a PostgreSQL database, equipped with the PostGIS module to manage geographical objects and accessed through the PyGreSQL module. It includes up to 45,000 among images, texts and cards obtained from reliefs in Mount Bego area. Each carved rock has a unique identifier number and precise GPS coordinates along with semi-structured annotations about the petroglyphs.

Three aspects form the basis of our work:

1. Integration and Semantic Annotation. Our main short-term goal is the integration of the Bicknell legacy, which is the most valuable source of knowledge on the petroglyphs that have been destroyed, into the Adevrepam database. Bicknell did not limit himself to faithfully depicting the petroglyphs, which would have been a remarkable contribution itself, but he also wrote semi-structured annotations, including location, rough description of the represented subjects and personal thoughts and interpretations. We remark that even if the Bicknell legacy is relatively small, a manual integration would be time-consuming and error-prone, besides being not scalable to our long-term goals. Therefore, we need to design a tool which helps domain experts integrate the data in a semi-automatic way. In this scenario duplicates are the major issue; the Bicknell legacy, in fact, contains drawings and annotations of petroglyphs which are already in the Adevrepam database. If this is the case, the integration tool needs to recognize duplicates in order to avoid the creation of two separate entries on the same subject. Recognizing duplicates is not straightforward, especially in the field of rock art, where different petroglyphs may share more or less the same patterns while being different. As far as semantic annotation is concerned, we aim at creating semantic relations among similar petroglyphs (that is those sharing the same patterns) thus allowing successive full and partial retrieval according to ontology-driven metadata and visual content.

2. Analysis and Normalization. The new integrated data must be checked against the existing ones for coherence. We recall that Bicknell was not an archaeologist and, even if the scientific community agrees on the importance of his work, some inaccuracies are still possible in his drawings. Moreover, to the best of our knowledge, an in-depth analysis and assessment of his work has not been done before. To this extent, we propose the exploitation of techniques based on both image similarity and ontology-driven metadata matching.

3. Interpretation and Semantic Enrichment. Since giving an interpretation to petroglyphs is not straightforward and two or more different interpretations may coexist and complement each other, we plan on developing a tool that helps domain experts to assess their interpretations and theories. Petroglyphs are elementary drawings, with shapes that can be found in many of them and so they are repeated. Given a shape, different interpretations can be as-

[1] http://www.rupestre.net/pdf_rtf/valca_bego_fr.pdf

signed to it. For example, a zigzag line is likely to represent water (de Lumley and Echassoux, 2009), a figure with two intersecting lines may represent a man and so on. Therefore, the system we are going to implement will be able to analyze all petroglyphs and, using ad-hoc sketch interpretation techniques, to take out all the "meaningful shapes" that occur frequently and to propose them to the domain expert. A foreseen usage scenario is that where the domain expert assigns an interpretation to each shape and adds ontology-driven metadata to that shape. Afterward, she can also access the petroglyphs to check whether its interpretation fits in the depicted scene, and the system provides suggestions of interpretations, thus being a real support for the experts. Moreover, the tool will also able to check for frequent co-occurrences of two shapes in the same petroglyph which may denote a possible joint interpretation of the two shapes (that is the two shapes have different interpretations if they occur in the same petroglyph).

In the following sections we will describe the technology we will use to implement these three steps.

3 Multi-agent System Architecture

As it is well-known MASs are an optimal solution when it comes to manage and organize data from multiple sources, which is the case in most Cultural Heritage applications such as ours. That is why we turned to a MAS as a basis of our framework.

The architecture of our MAS will extend the one discussed in (Casella et al., 2008) by adding specific classes of agents to it and enriching the system with ontology-driven knowledge, thus extending the communication interfaces with Semantic Web facilities. More specifically, the framework presented in (Casella et al., 2008) serves as a basis for a multi-domain sketch interpretation system, called *AgentSketch*, that recognizes and interprets symbols and simple shapes. *AgentSketch* uses either on-line or off-line interpretation of symbols; in the first case, the system starts the interpretation process while the symbol is being drawn, whereas in the latter the interpretation is carried out on a complete image. Sketch interpretation is the linchpin to implement our framework, as stated in Section 2 and *AgentSketch* works reasonably well on such tasks; therefore, we will use it while adapting it to our specific case. In the following we summarize the main features of Casella et al.'s agent framework (Section 3.1) and we show how we will extend it (Section 3.2). How we will use it is the subject of Section 5.

3.1 The Agent Framework

The main novelty of the agent framework in (Casella et al., 2008) over existing approaches is its flexibility, as it can be used in different contexts. Current solutions (Kaiser et al., 2004; Juchmes et al., 2005; Azar et al., 2006, just to cite some recent ones) either borrow techniques from stroke recognition, therefore limiting the set of symbols that can be recognized, or restrict themselves to particular domains, or they impose a-priori an usage mode (either on-line or off-line).

The agent framework (Fig. 1) is composed of four kinds of agents:

• *Interface Agent* (IA), that represents an interface between the agent-based framework and any input device used to draw a sketch. In the case of an online interpretation, this may be a pen-based device.

• *Input Pre-Processing Agent* (IPPA), that processes the input received from the Interface Agent and sends the obtained results to the Symbol Recognition Agents described below, using a format compliant with the interpretation approach they apply.

• *Symbol Recognition Agents* (SRAs), each one devoted to recognize a particular symbol of the domain by controlling one hand-drawn symbol recognizer (HDSR). HDSR are not considered agents since they are just passive providers of services within the MAS. They must be programmed by the system developer in order to recognize the symbols of the graphical language under consideration. SRAs may collaborate with other SRAs in order to apply context knowledge to the symbols they are recognizing, and with the Sketch Interpretation Agent described below that deals with the sketch interpretation activity. Of course, we must assume that a professional with skills on the computer science side specifies the patterns of symbols of interest, by defining one hand-drawn symbol recognizer for each of them. This activity is time consuming and expensive, but the domain expert involved in the activities described in this paper, namely the end users of the designed MAS, believe that it should be worth devoting time and resources to this activity, given the benefits that they could gain from it.

• *Sketch Interpretation Agent* (SIA) that provides the correct interpretation either of the sketch drawn so far (in case of an on-line drawing process) or of the entire sketch (in case of an off-line interpretation process) to the IA.

AgentSketch instantiates the general agent framework architecture and has been exploited for recognizing UML Use Case Diagrams in both on-line and off-line modes. It is implemented on top of Jade (Bel-

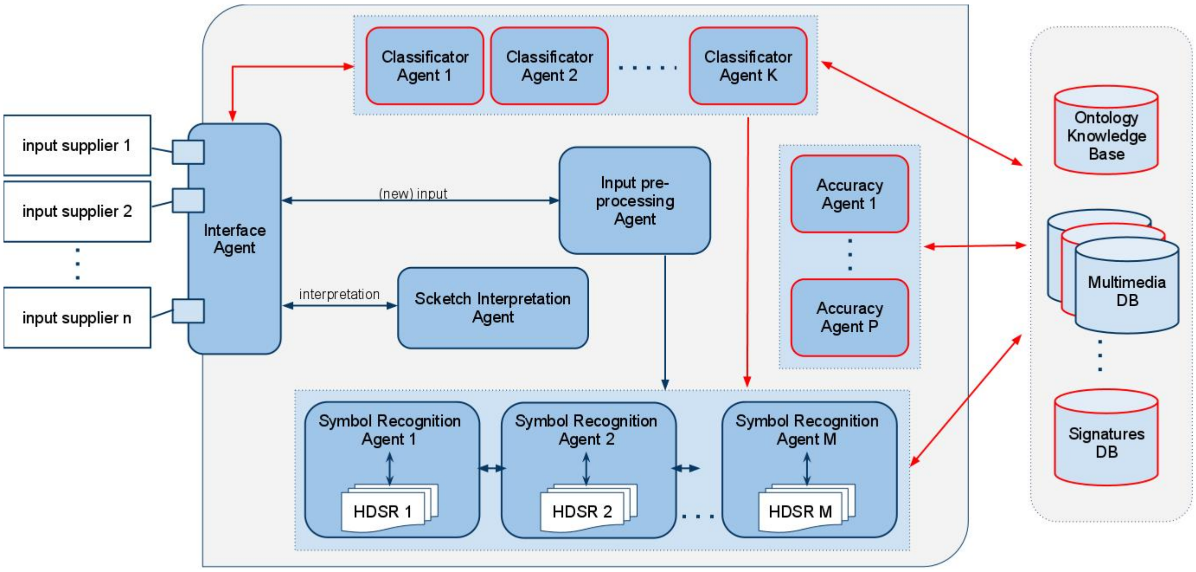

Figure 1: MAS architecture. Red components represent extensions to the original MAS architecture by Casella et al.. Red lines identify new communication interfaces that will be developed to allow agents interact.

lifemine et al., 2007).

3.2 Our extensions

We are working at the extension of Casella et al.'s architecture by adding agents with new functionalities and by defining ad-hoc communication interfaces among the agents and the data layer (namely the knowledge base, the multimedia DBs and the signatures DB in Figure 1). We have designed two agent types that will be added to the existing agent framework:

• *Classificator Agents* (CA), that will use both the interpretation of the sketch or image processed by the SIA and additional information that may be attached to the image itself (see Section 5.1) to classify the image (or sketch) according to the ontologies defining the MAS domain vocabulary.

• *Accuracy Agents* (AA), that will exploit CBIR techniques (Veltkamp and Tanase, 2002; Datta et al., 2008; Shishir K. Shandilya, 2010) by extracting a characterization (signature) of the images content (given ad-hoc heuristics, possibly formalized in the ontology), used for indexing and searching according to visual content similarity. This activity will be performed also for measuring the accuracy of Bicknell's drawings with respect to the more recent reliefs made by de Lumley's team.

The domain vocabulary of agents populating our MAS is provided by a set of ontologies (stored in the Ontology Knowledge Base, which we already started to fill) that may be both developed by hand, or extracted in a semi-automatic way from existing information sources (see Section 4).

4 Ontology Extraction

The MAS described in Section 3 heavily relies on the exploitation of ontologies for multimedia content classification and retrieval. We developed an ontology describing the petroglyphs found in Mount Bego (Papaleo et al., 2010, Figure 3). Our ontology was manually built and is based on the results of the archaeological reliefs of de Lumley and his team (de Lumley and Echassoux, 2009).

We remark that our aim is to integrate in our repository new data as they are available and this may require updating the ontology or creating new ones: it is self-evident that manual updates take time and are error-prone. Therefore we designed and implemented a *Role Ontology Extractor* tool that semi-automatically generates concepts and relations from texts, thus easing the process of creating new ontologies from scratch and/or updating existing ones. The *Role Ontology Extractor* tool extracts from a textual document the most relevant concepts that may be used in the MAS as well as relationships among them and their generalized super-concepts (Bozzano et al., 2010). The selection of meaningful relationships among the ones output by the extractor must be supervised by the domain expert: it is well known that Word Sense Disambiguation (Agirre and Edmonds,

2006) that the extractor applies in order to give each word its correct meaning within the context provided by the text, is an AI-complete problem, and a software tool can only support, but not substitute, the human user. Once the selection of meaningful relationships has been completed, the Role Ontology can be created from them: concepts are mapped into OWL *classes*, relationships among concepts are modeled as OWL *properties*, and the taxonomic relationships between concepts and their generalization is translated into the OWL *subClassOf relation*. The Role Ontology generated in this way can be used inside the MAS as a reference vocabulary among agents.

To show how the Role Ontology Extractor works, we run it on the following sentence from (de Lumley and Echassoux, 2009): "*These petroglyphs, [...] ,translate not just the daily preoccupations of these populations who needed rain, sources and lakes in order to fertilize their fields, but also their cosmological myths. At the center of these myths are the bull god, brandishing lightning, master of the storm and provider of fertilizing rain, and the high goddess, mother goddess or goddess earth, who needs to be fertilized herself by rain from the sky in order to bring abundance to humans.*"

The extracted meaningful relationships are listed below, and their ontological representation can be found in (Papaleo et al., 2010, Figure 4).

Quality → concept relationships

bull → god ; pastoral → population; agricultural → population; southern → alps; ancient → bronze age; daily → preoccupation; cosmological → myth; high → goddess; mother → goddess; goddess → earth.

Concept → action → concept

bull → brandish → lightning; god → brandish → lightning; population → need → rain; population → need → source; population → need → lake; source → fertilize → field; lake → fertilize →field; rain → bring → abundance.

The selection of *meaningful* generalizations and relationships from all those produced by the Role Ontology Extractor has been made by hand. Once this hand-made filtering stage has been completed, the OWL ontology can be generated in an automatic way. We are currently working at making the ontology compliant to CIDOC-CRM, a high-level ontology that enables information integration for Cultural Heritage data (Doerr, 2003) which is also known as standard ISO 21127:2006. The code of the Role Ontology Extractor, developed by Michele Bozzano as part of his Bachelor Thesis at the CS Department of Genoa University, is available under GPLv2 license (`http://www.disi.unige.it/person/MascardiV/Software/roleExtractor.html`). It was implemented using SWI Prolog extended with the ProNTo_Morph library for natural language processing (`http://www.ai.uga.edu/mc/pronto/Schlachter.pdf`).

5 The Framework at Work

This section shows how our framework will be used to integrate new data into the Adevrepam database and to retrieve content based on sketch interpretation.

5.1 Data integration

Classificator Agents (CAs) are responsible for integrating new data into the existing Adevrepam repository while updating the knowledge base and keeping the repository consistent. CAs work at the *semantic level* of images (photographs of Bicknell's drawings like those shown in the right hand side of Figures 3 and 4, new photographs or sketches of Mount Bego petroglyphs, as well as of petroglyphs dating back to the same period, but discovered in other places) and may take annotations from existing sources or specific metadata attached to images into account. A CA is activated when a new image is ready to be integrated. The CA sends the image to the Interface Agent (IA) responsible for the interpretation process. When the Sketch Interpretation Agent (SIA) completes the interpretation task, the IA sends the result back to the CA that triggered the interaction. This result is a set of concepts taken from the domain ontologies, stating the interpreted meaning of the set of recognized symbols, and used to tag the image. If metadata or annotations were attached to the image, the CA exploits them to provide an even more refined tagging, by identifying those concepts in the ontologies closest to the expected meaning of metadata. To carry out this task we developed and experimented with success an ontology matching algorithm that heavily exploits natural language processing (Mascardi et al., 2009).

If the set of concepts that tags the image is identical to the set that tags another image already stored in the database, the CA suggests to the user that the image might refer to content already present in the database and hence should not be considered as a new entry but as an update of existing information. On the other hand, if no similar tags are associated with any entry, the CA suggests that the image should be added as a new entry. Indeed, the CA must always interact with the domain expert in order to correctly add new content in the database. However, the CA provides suggestions to the human user that lighten the burden upon her.

After a decision has been made, the CA classifies the image according to the ontology concepts that tag it and stores it in the multimedia database with all the necessary metadata. Figure 2 illustrates the set of tasks and the communication among the agents in the MAS and the user.

After the integration step, the framework must ensure that the new data are correct and coherent with the existing data. In our specific case this means that we assess the accuracy of the drawings of Clarence Bicknell and to do so we need to compare his drawings with the data already in the repository. To this aim, the exploitation of techniques based on image similarity measures have been explored. An Accuracy Agent (AA) is activated by the user in order to establish if the selected image is part of one or more already existing classes of signatures (stored in the signatures DB) or if a new similarity class must be determined. If necessary, the new signature class is created and the signatures DB updated accordingly. Thus, inclusion, exclusion and intersection predicates are implemented to ensure partial similarity measures and ranking list of result answers. Once stored, the user can ask to an Accuracy Agent the similarity measure between two images. For example, it should make sense to measure the similarity between the images shown in Figure 3 and 4. Accuracy Agents operate at the *image level* by running algorithms that compute the similarity between images, without knowing nothing of their meaning.

Figure 3: Petroglyph identified by id. ZIIGIR3: de Lumley team's relief (left; private collection owned by Adevrepam) and Bicknell's relief (right; private collection owned by University of Genoa)

The assessment of Bicknell's drawings has never been done and is very important for the preservation of Mount Bego. Suppose that after this assessment we find out that Bicknell's drawings are accurate for every figure but the horned one; this would imply that we could digitally recreate all missing petroglyphs with high accuracy, except for those representing the horned figures.

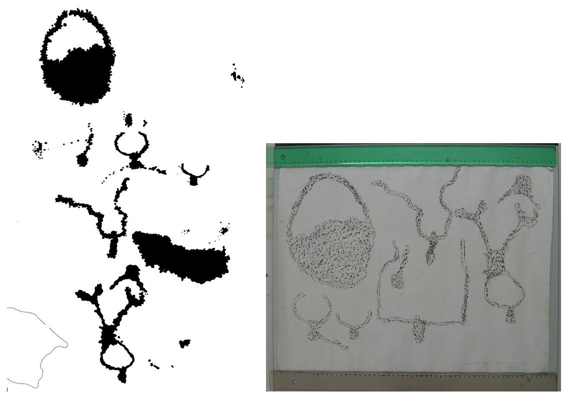

Figure 4: Petroglyph identified by id. ZVIIGIIR7: de Lumley team's relief (left; private collection owned by Adevrepam) and Bicknell's relief (right; private collection owned by University of Genoa)

5.2 Content-based Image Retrieval

Since Bicknell, the way archaeologists work today has dramatically changed. Instead of botanical paper sheets and pencil, they bring portable devices in their excavation campaigns, often equipped with sketch-based interfaces, as well as cameras and technical instruments for documenting with high precision their discoveries. Consider the scenario where an archaeologist making excavations either in the Mount Bego region or in any region where similar petroglyphs have been found (for example, the Valcamonica Valley[2]) discovers a new petroglyph. She may wonder if similar petroglyphs have been already recorded in the "Bronze Age petroglyphs" repository. To this aim, she may either take a picture of the petroglyph (Figure 6a) or draw a sketch by means of her PDA's sketch based interface (Figure 6b).

In both cases, the agents devoted to sketch interpretation may start their interpretation task (in on-line mode if the archaeologist sketched the petroglyph, in off-line mode otherwise) and discover that the pattern of the sketch (or of the picture given in input) respects the pattern of a known symbol, namely that of corniculates, based on the results of previous campaigns (Figure 5). The new sketches (or images) can be uploaded in the repository and structured according to the semantic data provided (ontology-driven and visual) thus enriching the repository with new multimedia content and new knowledge.

A Sketch Interpretation Agent can be activated also in the case in which a user wants simply to search for content in the repository "similar" to an input data. In this case, after the interpretation has taken place,

[2]http://www.rupestre.net/alps/valcamonica.html

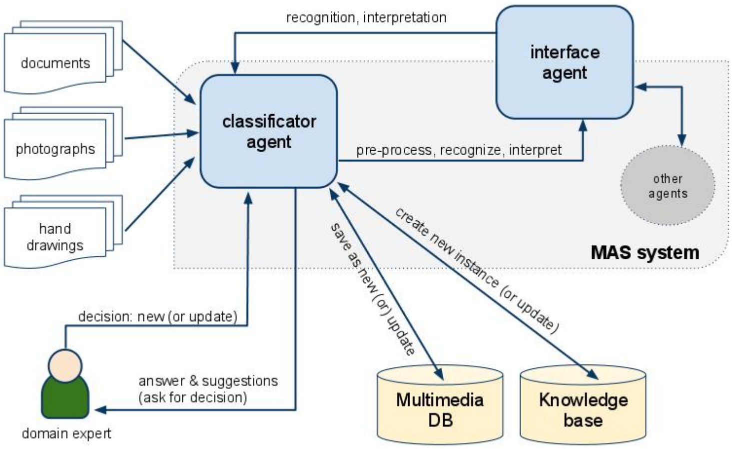

Figure 2: A schema describing the communication between a Classificator Agent and the Interface Agent in our MAS. The integration task will be performed under the supervision of the domain expert who, on the basis of the information provided by the Classificator Agent on tags, can decide whether the content must be added as "new item" in the multimedia DB or as an "update" of an existing entry.

a Classification Agent must compare the tags resulting from the interpretation process with those tagging content in the database, in the same way described in Section 5.1.

The Classification Agent will return a ranked list of multimedia content according to tag similarity criteria that can be different depending on users's needs (which can be defined using specific ontology-driven parameters). Also, users can decide to browse the repository according to the structure of the knowledge base, using specific ontology-driven paths and constraints, similarly to the approach presented in (Vrochidis et al., 2008).

Figure 6: Picture (from `http://www.cg06.fr/cms/annexes/merveilles/w_musee_merveilles/`) and sketch of an petroglyph.

6 Related and Future Work

As it is well known, a lot of research has been carried out the wide context of Cultural Heritage, leading to the development of many projects and tools. The most notable example is the "Epoch" Network of Excellence (`http://www.epoch-net.org/`, contract IST-2002-507382) recently concluded, which aimed at providing a clear organizational and disciplinary framework for increasing the effectiveness of work on the interface between technology and the cultural heritage of human experience. The Epoch NoE collected several tools for managing and organizing CH multimedia content, as for example AMA (Archive Mapper for Archaeology, `http://ama.ilbello.com/`), a web tool for mapping archaeological datasets to a CIDOC-CRM compliant format, or MAD (Managing Archaeological Data, `http://www.epoch-net.org/index.php?option=com_content&task=view&id=216&Itemid=332`) an application designed to store, manage and browse structured and unstructured archaeological datasets encoded in a semantic format. Other recent projects that shared similar goals with the Epoch NoE are Europeana, `http://www.europeana.eu/portal/` and 3D COFORM, Tools and expertise for 3D collection formation, `http://www.3d-coform.eu/`.

In (Vrochidis et al., 2008) a hybrid multimedia retrieval model is presented which provides a search engine capable to perform similar tasks to those we presented in this paper, without explicit use of a MAS.

Within this wide scenario, our key application is related to the preservation of Bronze age petroglyphs, by also taking advantage of the incredible valuable Bicknell's collection; we aim, as our ultimate goal, at designing a semantically annotated multimedia repository as a reference at European level as a thorough database of Bronze Age petroglyphs, which would be a definite contribution for domain experts, as rock art sites are spread all over Europe. To this pur-

Figure 5: Petroglyphs of corniculates of the Mont Bego region. From de Lumley, H. and Echassoux, A. (2009).

pose, we proposed a framework able to formalize the knowledge related to the available multimedia content through ontologies and to access and query the repository by using ad-hoc sketch interpretation algorithms. Besides the completion of the missing MAS components, their integration in the MAS and a careful testing, other future activities include, on the one hand to interact with the actors of Epoch in order to understand if our goal can be part of the more wide research community and on the other hand, to compare our approach with the one in (Vrochidis et al., 2008) in order to evaluate possible improvements. Finally, we point out that our framework focuses on the management, structuring and organization of multimedia content related to the Bronze Age but a lot of work could be done also in the "presentation" of this content to non-expert final users. We have already obtained preliminary results (Ancona et al., 2010) on interactions among 3D virtual worlds, living autonomous agents and semantic enriched multimedia content. The integration of those achievements on the presentation side within this framework will allow us to provide the domain experts, but also the potential virtual tourists and curious, with immersive experiences engaging on both the educational and ludic sides.

REFERENCES

Agirre, E. and Edmonds, P. (2006). *Word Sense Disambiguation: Algorithms and Applications (Text, Speech and Language Technology)*. Springer-Verlag.

Ancona, M., Mascardi, V., Quercini, G., Bogdanovych, A., de Lumley, H., Papaleo, L., Simoff, S., and Traverso, A. (2010). Virtual institutions for preserving and simulating the culture of Mount Bego's ancient people. In *VAST 2010*.

Azar, S., Couvreury, L., Delfosse, V., Jaspartz, B., and Boulanger, C. (2006). An Agent-based Multimodal Interface for Sketch Interpretation. In *MMSP-06*.

Bellifemine, F. L., Caire, G., and Greenwood, D. (2007). *Developing Multi-Agent Systems with JADE (Wiley Series in Agent Technology)*. John Wiley & Sons.

Bicknell, C. M. (1913). *A Guide to the Prehistoric Rock Engravings in the Italian Maritime Alps*. Tip. G. Bessone.

Bozzano, M., Locoro, A., Martelli, M., and Mascardi, V. (2010). Trattamento del linguaggio naturale tramite Prolog: un approccio promettente per generare istituzioni virtuali da testi scritti. In *CILC'10*. In Italian.

Casella, G., Deufemia, V., Mascardi, V., Costagliola, G., and Martelli, M. (2008). An Agent-based Framework for Sketched Symbol Interpretation. *J. Vis. Lang. Comput.*, 19(2):225–257.

Datta, R., Joshi, D., Li, J., and Wang, J. Z. (2008). Image Retrieval: Ideas, Influences, and Trends of the New Age. *ACM Computing Surveys*, 40(1).

de Lumley, H. and Echassoux, A. (2009). The Rock Carvings of the Chalcolithic and Ancient Bronze Age from the Mont Bego Area. The Cosmogonic Myths of the Early Metallurgic Settlers in the Southern Alps. *L'Anthropologie*, 113(5P2):969–1004.

Doerr, M. (2003). The CIDOC CRM An Ontological Approach to Semantic Interoperability of Metadata. *AI Magazine*, 24(3):75–92.

Juchmes, R., Leclercq, P., and Azar, S. (2005). A Freehand-sketch Environment for Architectural Design Supported by a Multi-agent System. *Computers & Graphics*, 29(6):905–915.

Kaiser, E., Demirdjian, D., Gruenstein, A., Li, X., Niekrasz, J., Wesson, M., and Kumar, S. (2004). A Multimodal Learning Interface for Sketch, Speak and Point Creation of a Schedule Chart. In *ICMI'04*, pages 329–330, New York, NY, USA. ACM.

Mascardi, V., Locoro, A., and Larosa, F. (2009). Exploiting Prolog and NLP techniques for matching ontologies and for repairing correspondences. In *CILC'09*.

Papaleo, L., Quercini, G., Mascardi, V., Ancona, M., Traverso, A., and de Lumley, H. (2010). Agents and ontologies for understanding and preserving the rock art of Mount Bego (extended version). Technical Report DISI-TR-10-11, Genova University, Italy. `http://www.disi.unige.it/person/MascardiV/Download/ICAART-tech-report.pdf`.

Shishir K. Shandilya, N. S. (2010). A Survey on: Content Based Image Retrieval Systems. *Int. J. of Computer Applications*, 4(2):22–26.

Veltkamp, R. C. and Tanase, M. (2002). Content-Based Image Retrieval Systems: A Survey. Technical report, Department of Computing Science, Utrecht University.

Vrochidis, S., Doulaverakis, C., Gounaris, A., Nidelkou, E., Makris, L., and Kompatsiaris, I. (2008). A hybrid Ontology and Visual-based Retrieval Model for Cultural Heritage Multimedia Collections. *Int. J. Metadata Semant. Ontologies*, 3(3):167–182.