International Journal of Innovative
Computing, Information and Control
Volume **XXX**, Number **0XXX**, January **2014**

ICIC International ©2011 ISSN 1349-4198
pp. xxx

A HOLONIC MULTI-AGENT SYSTEM FOR SKETCH, IMAGE AND TEXT INTERPRETATION IN THE ROCK ART DOMAIN*

Viviana Mascardi[1], Daniela Briola[1], Angela Locoro[1], Daniele Grignani[1], Vincenzo Deufemia[2], Luca Paolino[2], Nicoletta Bianchi[3], Henry de Lumley[3], Damiano Malafronte[1], Alessandro Ricciarelli[1]

[1] Università di Genova, Via Dodecaneso 35, 16146, Genova, ITALY
viviana.mascardi@unige.it, daniela.briola@unige.it, angela.locoro@unige.it, grignani.d@gmail.com, damianomal@gmail.com, ricciarelli.alessandro@gmail.com

[2] Università di Salerno, Via Ponte don Melillo, 84084 Fisciano (SA), ITALY
deufemia@unisa.it, lpaolino@unisa.it

[3] Lab. Dép. de Préhistoire du Lazaret, 33 Bis Boulevard Franck Pilatte, 06300 Nice, FRANCE
nbianchi@lazaret.unice.fr, lazaret@lazaret.unice.fr

Abstract. This paper presents the architecture of a holonic multi-agent system for rock art interpretation and discusses the results achieved within the "Indiana MAS" project. We show how the AgentSketch and the ImageRec holons belonging to the Indiana MAS, able to cope with hand drawn sketches and images respectively, have been tested in the domain of Mount Bego's prehistoric rock art (southern French Alps), and how the Manent agent-based framework for the seamless integration of Digital Libraries has been plugged into Indiana MAS to provide text classification, as well as multilingual access to structured repositories. The way Indiana MAS holons cooperate in order to provide correct interpretations of ambiguous shapes is discussed by means of an example based on hypotheses recently advanced by archaeologists.

1. Introduction and Motivations.

The "Indiana MAS and the Digital Preservation of Rock Carvings: A Multi-Agent System for Drawing and Natural Language Understanding Aimed at Preserving Rock Carvings" project ("Indiana MAS" for short[1]), funded by the Italian Ministry for Education, University and Research, MIUR, and spanning from March 2012 to February 2015, aims at developing a technology platform based on intelligent software agents for the digital preservation of rock carvings, which both integrates and complements the techniques usually adopted to preserve heritage sites. The platform will support domain experts in the creation of a repository, which may become a reference at Italian and European level as a thorough database of rock art, and in the interpretation of rock carvings. It will also promote the awareness and the preservation of the cultural treasure by making cultural information accessible to everyone on the Internet, and by preserving it in a digital format, for future generations. To this end, the Indiana MAS platform will enable the preservation of all kinds of available data about rock carvings, such as images, geographical objects, textual descriptions of the represented subjects. It will provide the means to organize and structure such data in a standard way and will supply domain experts with facilities for issuing complex queries on the data repositories and making assumptions about the way of life of the ancient people.

* *This article is a revised and extended version of a paper presented at the 6th International KES Conference on Agents and Multi-agent Systems – Technologies and Applications 2012* [42].

[1] We use "Indiana MAS" to denote both the funded project and the Multi-Agent System, or MAS for short, that will result from it.

The choice of agent technology for addressing the Indiana MAS goals was a very natural one, given the need that each component of the system, while operating in a highly autonomous way, interacts and coordinates with the other components to share information and to reason about it in the most effective way. However, designing Indiana MAS as a "flat" MAS could not cope with all of its challenging goals. For this reason we opted for a more flexible approach based on holons: the three key services offered by Indiana MAS (sketch recognition, image recognition, and multilingual access to digital libraries) are provided by systems that are themselves MASs, and that are seen as black boxes by the Indiana MAS agents.

In this paper we present the three main components of Indiana MAS (namely, AgentSketch, ImageRec, and Manent) and the results of the experiments performed during the first nine months of the research project. The tests we conducted until now consider each component as a separate unit since their full integration is still under way.

The paper is organized in the following way: Section 2 surveys the state-of-the-art on holonic MASs, ontologies, sketch and image recognition, and digital libraries, with a special attention to agent-based approaches. Section 3 introduces the Indiana MAS architecture and functionalities. Sections 3.2.1 and 3.2.2 describe the AgentSketch and ImageRec holons, respectively, and discuss how they have been tested in the domain of Mount Bego's rock art. Section 3.2.3 discusses the Manent holon for classifying texts based on the Indiana Ontology, collecting information coming from distributed Digital Libraries, providing multilingual services, and discusses its integration within Indiana MAS. Section 4 illustrates how the correct interpretation of images and sketches can result from the interaction between the holons and the other Indiana MAS components. It also provides at the same time a concrete application of our study, as well as a practical example of usage of the system. Section 5 surveys related projects and outlines the future developments of our work.

2. **State-of-the-Art.**

2.1. **Holonic Multi-Agent Systems.**

A *holarchy* is a connection between *holons* – where a holon is both a part and a whole. The term was coined by A. Koestler in his book *The Ghost in the Machine* published in 1967 and derives from the Greek word *holos*, meaning whole, and *on*, meaning part. Each holon exists simultaneously as both a distinct entity built from a collection of subordinates and as part of a larger entity. According to [32], this organizational paradigm is suitable to design MASs

> *where goals can be recursively decomposed into subtasks that can be assigned to individual holons. Given such a decomposition, or a capability map of the population, the benefits the holonic organizations provide are derived primarily from the partially autonomous and encapsulated nature of holons.*

The industrial application domain where holonic MASs have been most widely exploited since the beginning is manufacturing [19, 24, 33, 58], but real-time control, supply chain management, resource allocation, production planning/scheduling, air traffic, and e-health are suitable domains as well, as witnessed by the HoloMAS Conference Serie on Industrial Applications of Holonic and Multi-Agent Systems that reached its sixth edition in 2013 (`http://cyber.felk.cvut.cz/HoloMAS/2013/`).

Due to the ever growing importance of this paradigm, ad hoc methodologies for modeling and developing holonic MASs have been recently proposed inside the MAS community [13].

2.2. **Ontologies.** Semantic web applications and ontology based systems are growing all around the world: this is due to many factors, above all the need to have a formal representation of knowledge that can be shared among computers and humans. An ontology is a formal way to represent concepts and data, or

> *[...] a specification of a conceptualization.* [28].

Ontologies have become the cornerstone for a standard, formalized, and shared semantic description of information systems and of their actors, as well as the main chance to achieve interoperability and integration in common environments. Different domains may be represented with an ontology, from taxonomies to complex system's models with axioms and rules. An ontology conceptualizes a domain of discourse through the description of its entities and their relationships, and the rules and constraints inside the domain itself that ought to be explicitly stated.

Ontologies are widely used inside MASs since they represent a concrete means to achieve semantic interoperability among agents: managing one or more ontologies in a MAS allows unknown agents to communicate and understand each other [29]. This close relationship between ontologies and MASs is proven by a huge literature on the topic: many existing MASs exploit ontologies to reach theirs goals or to formally represent their knowledge [4, 12, 45]. Furthermore, MASs are quickly moving towards the open world of Internet, that is rapidly becoming the Semantic Web also thanks to the adoption of the ontologies: hence the natural cross-fertilization between the two technologies [15, 35, 43, 55].

2.3. **Sketch and Image Recognition.** In the recent years there has been a growing interest in sketch recognition, partly due to the increasing popularity of pen-based input devices. However, the free drawing style of sketching makes the construction of robust sketch recognition systems a very difficult task. Indeed, the lack of precision during the sketching process and the potential variability in the way users draw shapes (even if they are drawn by the same user) make the interpretation of the sketched symbols harder. To improve the recognition performance many approaches exploit contextual information to resolve pending recognition ambiguities, such as SketchREAD [1] and Tahuti [31], or to recover from low-level interpretation errors [14] or, in combination with visual appearance features, to make the recognition less sensitive to noise and drawing variations [44].

The literature describes many proposals where the interpretation of sketches is performed by applying the same recognition approach to any symbol in a sketch. Allowing the adoption of different techniques for recognizing different symbols, as we do in AgentSketch and ImageRec, is a powerful approach for coping with the heterogeneity of symbols and the need to make the system as flexible and scalable as possible. Thus, traditional techniques are more frequently being substituted by new systems based on agents, where each agent is specialized in recognizing a subset of symbols or patterns in images, and the final interpretation of either the individual symbol or the complete scene results from interaction among agents.

Mackenzie and Alechina [41] use agents for the classification and understanding of childlike sketches of animals, using a live pen-based input device. Juchmes *et al.* [34] based their freehand-sketch environment for architectural design on a MAS. Their system interprets the strokes drawn on the screen thanks to the activity and collaboration of different types of agents. Azar *et al.* [3] extend this architecture with the possibility of interpreting vocal information. More recent agent-based approaches for sketch recognition have been presented by Fernández-Pacheco *et al.* [23] and Avola *et al.* [2]. The first introduces a system organized on two levels of agents for supporting the recognition process: primitive, lower level agents are in charge of the syntactic recognition, whereas complex, higher level agents exploit contextual information to carry out the semantic recognition. The latter employs two types of agents for sketch recognition: feature evaluators, which compute a specific feature on the drawn sketch, and symbol recognizers, which read from the pool of features those needed to discriminate the symbol for which they are responsible. A mediator agent is applied in case of conflicts.

As far as the recognition of objects within images is concerned, the literature is huge. The difficulty of object detection is due to the fact that objects have complex appearance patterns and spatial deformations. Various approaches have been used to represent objects in the detection process, such as color information, texture, edge orientation, and so on. In our experiments we consider Haar-like features and SIFT descriptors for the detection of petroglyphs within colored images. The first has been widely used in the computer vision community starting from their use in the Viola and Jones' real-time object detection [54], while SIFT has been shown to be an effective descriptor for traditional object recognition applications in static images [40].

For analyzing the similarity between petroglyph reliefs in black and white images we consider two widely used descriptors for shape recognition: *Shape Context* [5] and *Radon transform* [46]. The first is based on the observation that a set of vectors connecting a point on the shape to the rest of the shape points constitutes a rich and discriminative description of the shape. The *Radon transform* has several properties that make it appealing for shape matching: it is a lossless transform, has low complexity, and has useful properties concerning rotation, scaling, and translation transformations.

2.4. **Digital Libraries.** A Digital Library is

> *A potentially virtual organization, that comprehensively collects, manages and preserves for the long depth of time rich digital content, and offers to its targeted user communities specialized functionality on that content, of defined quality and according to comprehensive codified policies.* [8]

As explained by its definition, the term Digital Library may encompass heterogeneous systems that range from the simple digitization of material libraries and analog archives to the definition of connected collective information spaces in which humanity access, exchange and produce new knowledge in a friendly, multi-modal, multi-cultural and multi-language way. As a consequence, a Digital Library system may contain simple digital objects with their metadata, or complex and advanced service infrastructures for the description, the storage, the retrieval, and the processing of digital information and knowledge [53]. These advanced technologies rely on scientific outcomes from disparate research strands: from data management to artificial intelligence, from semantic web to human-computer interaction, and so on. The core entities of a digital archive are the collections of information objects structured as contents with annotations and metadata. In particular, metadata have the main task of providing the syntactical, semantic and contextual interpretation of an information source. For this reason, they are important for preserving the interoperability of systems and then the access to the resources from whatever device in the world. A standard protocol that Digital Libraries exploit for exposing their own archives and harvesting worldwide repositories is the Open Archives Initiative Protocol for Metadata Harvesting, OAI-PMH [18], an application-independent interoperability framework where a client-server architecture is built on top of *data provider* and *service provider* roles. The metadata exchanged through the protocol are, for example, in Dublin Core format [16], for the description of single digital objects. As for their complexity and multi-service structure, Digital Libraries often rely on agent-based architectures. In [56] a mix of both stationary as well as mobile agents are applied to data management tasks such as parallel queries and integration of data from different digital repositories. In [48] interface agents are designed to facilitate users to cope with the complexity of digital libraries information systems, and an agent-enabled digital library architecture, called AGS, has been conceived for making agents available as a service to which navigational, direct manipulation and data management tasks can be delegated. Derbyshire et al. [20] argue that agent-based digital libraries are able to drive information economies seen as pro-actively self-adapting framework where "computational agents carry out the roles of information producer, consumer, and broker".

3. Indiana MAS Architecture and Functionalities.

3.1. The Indiana MAS High-Level Architecture. Indiana MAS [39, 22] integrates intelligent software agents, ontologies, multilingual natural language processing, sketch and image recognition techniques. Ontologies allow to define a common vocabulary that can be profitably exploited to organize data associated with rock carvings, included their semantic annotations, and to create semantic relationships between them. Natural Language Processing (NLP) techniques are used for extracting relevant concepts from textual documents in order to categorize them and for mining semantic relationships among them, hence supporting both digital objects classification and ontology evolution. Multilingual issues arise because texts we consider in our project may be written in English, French, and Italian. Finally, sketch and image recognition techniques are applied to classify the elementary shapes of the carving drawings and pictures, and to associate their possible interpretations with them.

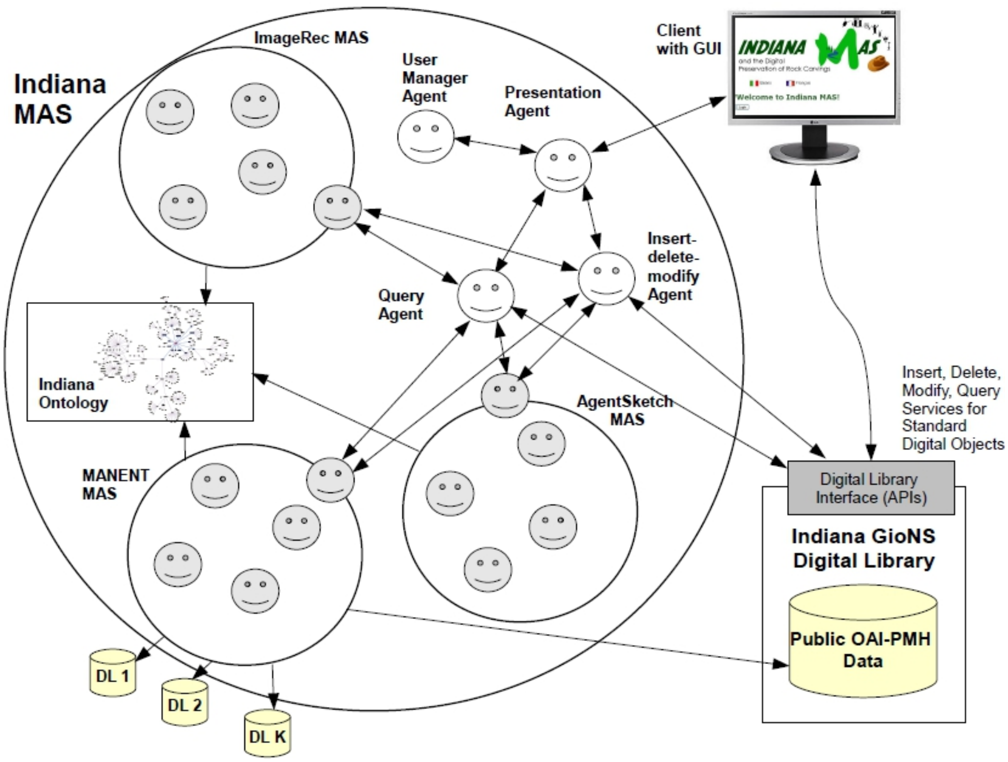

FIGURE 1. Indiana MAS architecture. Only intra-agents links are shown; some arrows representing access to the Indiana Ontology are omitted for sake of readability.

Figure 1 describes the Indiana MAS holonic architecture whose main components are holons, that is, the designed and developed MASs[2] are at the same time "part" of a

[2]ImageRec has been specifically designed for the Indiana MAS project, whereas AgentSketch and Manent were conceived before the project's start, and have been extended for the project's purposes.

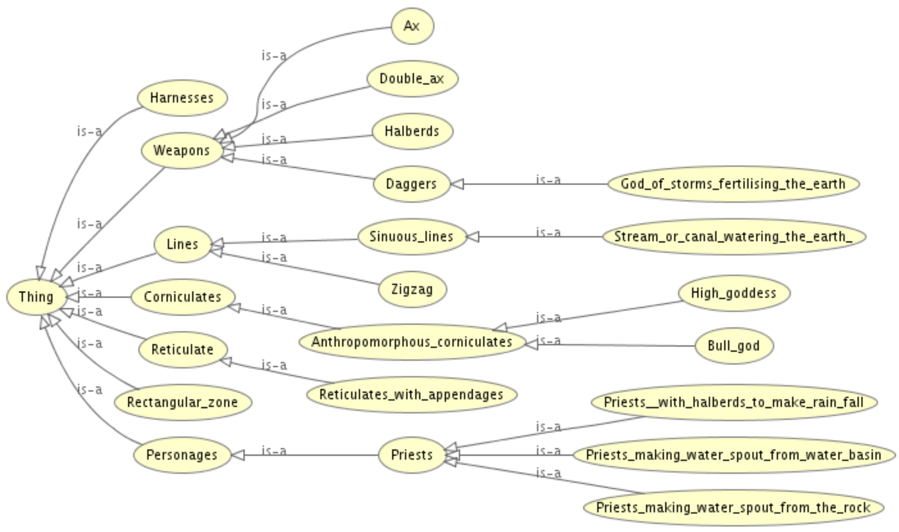

FIGURE 2. Part of the Indiana Ontology: known petroglyph.

bigger MAS (i.e., the Indiana MAS, the higher "whole") that contains them, as well as independent MASs (i.e., self-contained and autonomous "wholes", recursively containing their "parts") [27].

The main components of Indiana MAS are:

- The Indiana Ontology, partly shown in Figure 2, that structures the domain of interest of the project, namely that of Mount Bego Rock Art.
- AgentSketch holon, described in Section 3.2.1, that interprets manual drawings, based on the Indiana Ontology.
- ImageRec holon, described in Section 3.2.2, that performs image recognition and classification based on the Indiana Ontology.
- Manent holon, described in Section 3.2.3, whose main purpose is accessing Digital Libraries spread around the world, besides the Indiana GioNS one, and to classify multilingual documents available there by applying NLP techniques both to the document textual metadata and to the text body. Manent also provides to Indiana MAS a service for text classification based on the Indiana Ontology.
- The Indiana GioNS Digital Library, which will contain all the Digital Objects inserted into the system by registered users, together with their metadata, needed for their later retrieval; it is kept outside of Indiana MAS to mean that it is a traditional Digital Library that can be accessed from the outside, by using the suitable standard protocols (for example the OAI-PMH protocol), while it will expose some APIs to let Indiana MAS manage it. The Indiana GioNS Digital Library will be filled with structured and unstructured multilingual and multimedia digital objects related to Mount Bego's rock art, coming from the Bicknell Legacy owned by the University of Genoa[3], the ADEVREPAM database owned by the Laboratoire Départementale

[3]The explorer who first realized the importance of Mount Bego carvings from an historical point 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 [6]. Bicknell's collection includes inedited drawings owned by the University of Genoa that amount to 16,000 tracings on different materials, annotated with notes in English and in Italian, and partly digitized.

de Préhistoire du Lazaret[4], and other documents by archaeologists working on sites similar to Mount Bego's one. Every digital object inserted into the Indiana GioNS library is characterized by standard metadata, such as *title* and *author*. We will identify a set of standard metadata that will contain domain-specific values (that will be called "domain-specific metadata" in the next): these metadata and their values may vary depending on the digital object they refer to, and all the domain-specific values are chosen from the Indiana Ontology. Simple digital objects can be linked to form more complex objects.

- Client with GUI, for interacting with Indiana MAS (that is, to interact with the Indiana GioNS Digital Library in a transparent way to the end user), that exploits the over mentioned APIs: the GUI allows the user to access the Digital Library to insert, retrieve, and modify digital objects (with respect to the user's permissions) using metadata, as for a standard Digital library, and also provides the functionalities that are specific for the Indiana MAS domain, that include query by sketch, query by image similarity, query by text similarity. The GUI will help the user create complex queries over metadata, supporting the user in the selection of the possible values for domain-specific metadata and in the creation of the query considering the relationships among digital objects. The GUI will also let Indiana MAS administrators manage users and their permissions.
- Presentation Agent, for suitably presenting the content to the users depending on their profiles, permissions and device.
- User Manager Agent, for personalization and profiling purposes.
- Insert-Modify-Delete Manager and Query Manager Agents, for preparing the right queries/requests to the Indiana GioNS Digital Library. The Indiana GioNS Digital Library can of course be queried in a traditional way using the standard protocols that it offers to the world, but this would allow users to make queries that are strictly limited to metadata, without taking advantage of the Indiana MAS full potential. In particular, if - for example - a user performs a "query by sketch of a single symbol", this query must be split into different tasks that can be offered only by Indiana MAS, involving:
  1. sending the sketch to AgentSketch for an interpretation and getting back the interpretation (list of terms from the Indiana Ontology) for that symbol, of the form {[corniculate], [priest],...};
  2. preparing a query to the Indiana GioNS Digital Library, using the APIs it supports, that asks "Find all the Digital Objects that are sketches of individual symbols, and that are classified either as corniculates or as priests";
  3. sending the results back to the Presentation Agent, filtered according to some criteria (if any).

Indiana MAS can be accessed by registered users, that have high access privileges and can modify the Indiana GioNS Digital Library objects, and guests that are only allowed to query the information sources managed by Indiana MAS. Both kinds of users access Indiana MAS by means of a User Manager designed along the lines of old and well-known "digital butlers" [11, 36], providing user profiling and content personalization capabilities, as well as collaborative filtering and search.

[4]Many years after Bicknell's campaigns, several teams led by Henry de Lumley have been surveying and mapping the Mount Bego area starting from 1967. The ADEVREPAM database stores images of the tracings made by De Lumley's team with metadata in English and French.

The management of queries is demanded to the Query Manager that sends requests to the holons, receives and integrates all the answers, and sends the integrated results back the the Presentation Agent.

In order to make the management of users' requests efficient, we designed a control flow based on a lazy creation of agents and holons performed when existing instances - if any - are too busy to efficiently cope with a new request. This approach will allow Indiana MAS to be highly scalable still avoiding both bottlenecks and useless replication of agents.

3.2. **The Indiana MAS Components: Overview and Experiments.**

3.2.1. *The AgentSketch MAS.* AgentSketch [9] consists of the agents shown in Figure 3.

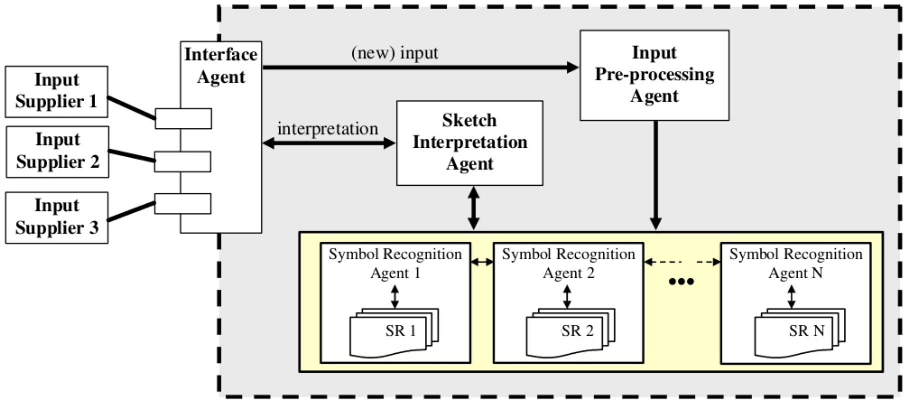

FIGURE 3. AgentSketch architecture

The Interface Agent represents an interface between the agent-based AgentSketch framework and the generic Input Suppliers that are not included inside the framework. The nature of these input suppliers may vary according to the type of sketch to be interpreted. The Interface Agent converts the information produced by the input suppliers into a suitable format for the Input Pre-Processing Agent (that, in turns, contacts the Symbol Recognition Agents); it also interacts with the Sketch Interpretation Agent (SIA) for sending the sketch interpretation requests to it, and for delivering its answers to the user.

The Input Pre-Processing Agent processes the input received from the Interface Agent and sends the obtained results to the Symbol Recognition Agents described in the following, using a format compliant with the recognition approach they apply.

Each Symbol Recognition Agent (SRA) is devoted to recognize a particular symbol of the domain. Moreover SRAs may collaborate with other SRAs in order to apply context knowledge to the symbols they are recognizing, and with the SIA that deals with the sketch interpretation activity. On-line recognition of hand-drawn sketches is demanded to Symbol Recognizers (SRs in the figure) that are not agents, lacking most of the agent-characterizing features, but just software modules managed by SRAs. As long as there is one SRA that correctly integrates SRs by managing their execution as well as data conversion issues, the actual implementation of the SRs and the approach to recognition that they adopt do not matter. The inherently flexible and modular agent-based approach allowed us to seamlessly cope with symbols that have been recognized by heterogeneous SRs managed by ad-hoc SRAs.

The Sketch Interpretation Agent provides the correct interpretation of the sketch drawn so far to the Interface Agent. In particular, it analyzes the information received from SRAs

```
(define shape Corniculate
  (description
    “Symbol representing a corniculate in Mount Bego’s carvings”
  )
  (components
    (Ellipse head) (Line horn1_lower) (Line horn2_lower) (Line horn1_upper) (Line horn2_upper)
  )
  (constraints
    (touches head horn1_lower) (touches head horn2_lower)
    (hornsConnected horn1_lower horn1_upper) (hornsConnected horn2_lower horn2_upper)
    (equalLength horn1_lower horn2_lower) (equalLength horn1_upper horn2_upper)
    (equalLength horn1_lower horn2_lower)
    (over horn1_lower head) (over horn2_lower head) (over horn1_upper horn1_lower) (over horn2_upper
horn2_lower)
  )
)
```

Figure 4. LADDER description of a corniculate.

and solves conflicts between symbols that might arise, taking information available into the Indiana Ontology into account as shown in Section 4. When all the conflicts have been solved, the SIA proposes the sketch interpretation to the Query Manager that delivers it to the Presentation Agent.

AgentSketch has been experimented in the domain of Use Case Diagrams [9]. Its exploitation for reasoning about hand-drawn sketches in the physical security domain has been discussed in [10]. In order to recognize rock art carvings while they are being sketched, we implemented a SRA working on-line and based on LADDER [30]. In LADDER, symbol recognition is performed using the rule-based system Jess [26]. In particular, for each symbol of the domain, a Jess rule is automatically generated from a LADDER structural shape description, which mainly contains information on the shape of the symbol. Recognition using Jess is sensible neither to the order of the strokes, nor to the symbol dimension. These features make the approach very stable.

As an example, the LADDER description of a corniculate is shown in Figure 4. This rule defines a corniculate as one ellipse and four lines from which the shape is built, plus the topological constraints defining the relationships among these elements: two lines (representing the lower part of the horns) must be similar in length, touch the ellipse (representing the head), and must be over it; the two other lines (representing the upper part of the horns) must be similar in length, and each of them must be above the lower part of one horn and connected to it.

The final behavior of this rule is that, when in the working memory of the LADDER application there are four lines and one ellipse that respect the precondition of the rule, the rule is fired and a corniculate symbol is recognized. A similar (but simpler) rule has been defined for recognizing short weapons.

The shapes described with LADDER must be diagrammatic or iconic since they have to be drawn using a predefined set of primitive shapes and composed using a predefined set of constraints.

Figure 5 shows a screenshot of AgentSketch at work. Figure 5(a) shows the symbols sketched by the user, and 5(b) shows which of them have been correctly recognized as corniculates (black symbols) and as daggers (red – or light gray if printed in gray-scale – symbols). In this example, all the symbols have been correctly recognized as belonging to their own category, apart from the dagger in the lower left portion of the screen.

Table 1 shows the results of 220 experiments. Each rock-art element was drawn 20 times, 10 of which in a correct way, and 10 in a wrong way (correct and wrong w.r.t. the

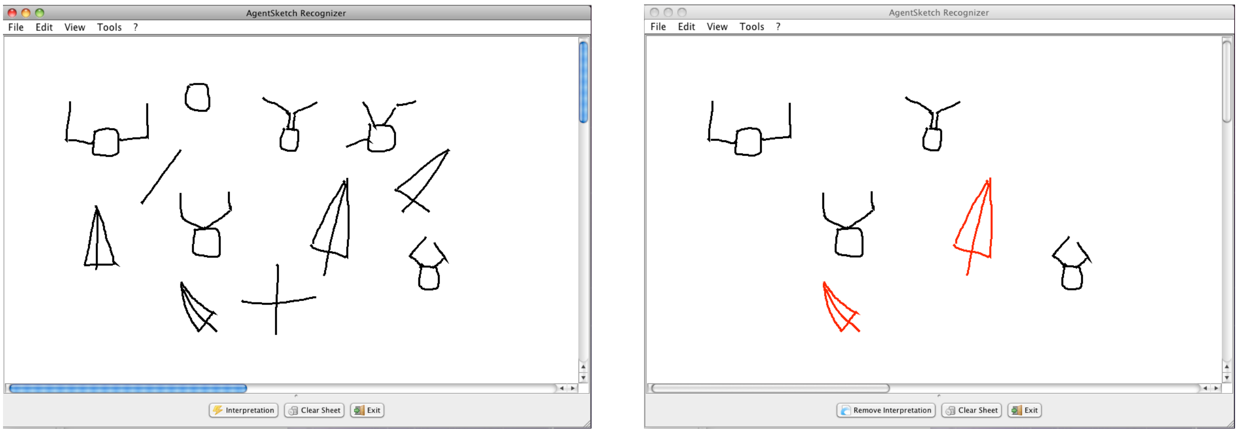

a b

FIGURE 5. AgentSketch GUI: a user's sketch (a) and the recognized symbols (b).

shape that AgentSketch should recognize). In order to evaluate the results, we used the following standard information retrieval indicators:

- True positive (TP): the symbol was correctly identified
- False positive (FP): the symbol was incorrectly identified
- True negative (TN): the symbol was correctly rejected
- False negative (FN): the symbol was incorrectly rejected
- Sensibility or True Positive Ratio (TPR): TP / (TP + FN)
- False Positive Ratio (FPR): FP / (FP + TN)
- Accuracy (ACC): (TP + TN) / (TP + TN + FP + FN)
- Specificity or True Negative Ratio (TNR): 1 - FPR

Symbol	**TPR**	**FPR**	**ACC**	**TNR**
Halberd	75%	5%	85%	95%
Axe	100%	0%	100%	100%
Corniculate	90%	7,5%	90%	92,5%
Dagger	75%	0%	88%	100%

TABLE 1. Results of our experiments on rock art symbols

As we can see from Figure 5, despite the simplicity of their shape, correctly recognizing daggers raises some problems. Daggers that are sketched almost in the same way may be either recognized or not, depending on very small variations of the length and relative position of their constituent elements. Relaxing the constraints of the rule for recognizing daggers would eliminate this problem, but would lead to detection of false positives that, in the current setting, are completely absent. Similar considerations hold for halberds.

3.2.2. *The ImageRec MAS.* The architecture of the ImageRec holon is similar to the AgentSketch's one: agents devoted to recognize specific symbols in images are coordinated by an higher-level agent, the Image Interpretation Agent (IIA), that can integrate the partial information obtained by SRAs, to suggest an interpretation of both the individual elements, and the full image. In this setting, the IIA may ask more SRAs, operating using different algorithms, to classify objects within a given image according to the Indiana Ontology, and then may accept the classification of the "most reliable" SRA. Reliability may be defined at design time (for example, as discussed later in this section we know that SRAs using Radon on black and white images recognize oxen better than SRAs using Shape Context, and hence they will always be considered more reliable than SC-based SRAs when oxen are involved) or at run-time, implementing some trust-based mechanism that also involves the user in the loop.

The main difference w.r.t. AgentSketch is the kind of objects (images instead of hand-drawn sketches) ImageRec operates on, that leads to the exploitation of completely different recognition techniques. The SRA for recognizing rock carvings in colored images has been developed using an existing open source library of programming functions for real time computer vision, OpenCV (Open Source Computer Vision[5]). OpenCV is widely used for object detection as well. Among the many functions offered by OpenCV, we used those that exploit Haar-like features [54] for categorizing subsections of images based on the intensity of pixels in the region, and the AdaBoost machine learning algorithm [25] for training classifiers to recognize objects in images given positive and negative samples.

The implementation of another SRA based on SIFT [40] is under way. SIFT is a method for detecting and describing local features in images and has been recently used for automatic coin classification [57]. It could turn out to be a successful choice since rock engravings and ancient coins show many similarities from a computer vision viewpoint. In the coin case, SIFT could recognize the same object/element/figure in different coins even if the actual shape varies from one item to another.

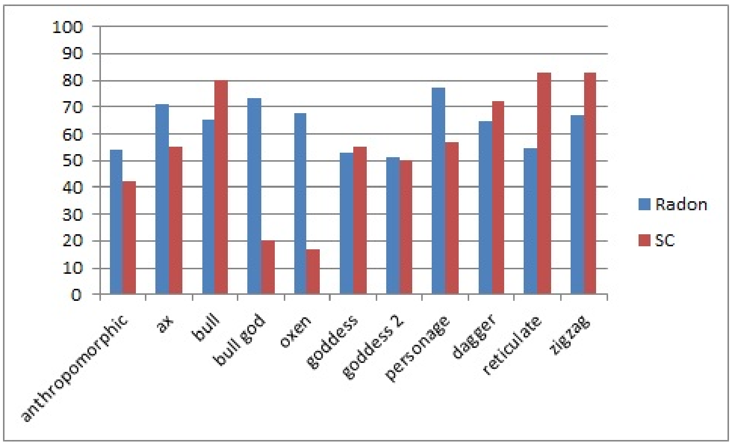

FIGURE 6. Percentage of correct classifications considering the top scored class using Radon and Shape Context.

Unfortunately, as documented in [50, 59], blurriness, sketchiness, and incompleteness make the automatic, as well as the human, interpretation of rock carvings a very difficult issue that requires a deep research even for identifying where carvings are onto the stones. For this reason, we decided to also focus on black reliefs on a white background at least in the first stages of the project.

To provide a similarity measure between petroglyph reliefs we evaluated two different descriptors, namely: Shape Context [5] and Radon transform [46], which allow to measure the similarity of petroglyph images on the basis of their shapes. However, since the use of these approaches is not fully suitable for the Web context where almost-real-time answers are required, we combined these approaches with Self-Organizing Map (SOM) [21], a clustering technique based on neural networks. Generally, the combined approach may be summarized into three steps: 1) extracting the descriptors from the training images and associate them with the petroglyph classes, 2) feeding the SOM with the descriptors,

[5] `http://opencv.willowgarage.com/wiki/`. Last accessed: February the 13th, 2013.

3) querying the SOM by searching for the cluster with the less distance from the query image descriptor.

To perform an objective comparison between the two approaches and to evaluate which one has the best performance in terms of effectiveness we carried out a preliminary experiment on a dataset composed of 1400 relief images. In particular, the dataset contains 1400 images, divided into 11 classes, which have been computed by means of an image distortion algorithm applied on a set of 74 petroglyph images from Mont Bego. The dataset has been processed during the experiments by using a modified version of the k-fold approach [47], with $k = 5$. In particular, for k times a set of $m = 11$ images − one for each class − have been chosen from the 74 original images, and the corpus of $N = 1406$ images has been split into two subsets: a subset of $[N - (m * 19)] = 1197$ images was used for training, while the remaining $m * 19 = 209$ images were left out for testing. In particular, at each iteration, we extracted for testing m images from the dataset together with their deformed version.

Figure 6 summarizes the recognition rates we obtained from the experiments. Basically, each bar represents the rate between the number of top-scored correct classifications and the number of queries for each class. It is worth to note that no algorithm generally provides optimal classification for all the petroglyph classes. As we can notice, Shape Context works better for *zigzag* and *reticulate* classes, but worst respect to Radon when considering the *oxen* and *personage* classes. For this reason, it is important to have two or more classifiers which work together in order to have always the best performance in terms of effectiveness.

3.2.3. *The Manent MAS.* Manent [37, 38] is a system composed of a metadata harvester, as well as of services capable of providing multi-language text processing, and content classification based on the Indiana Ontology. In Manent, some components are conceived

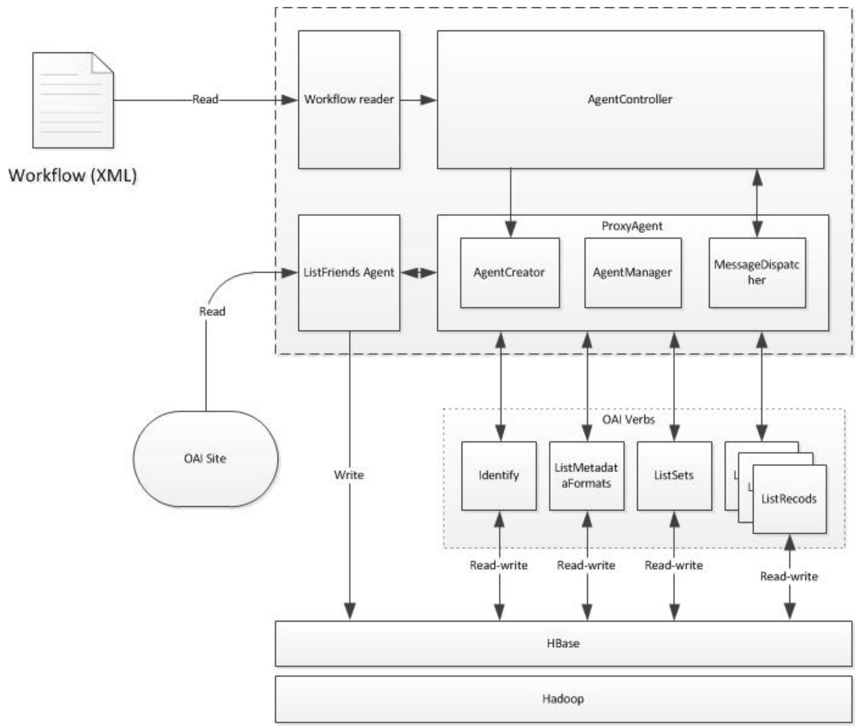

FIGURE 7. The Manent MAS harvester.

to work as agent-based services. The Manent OAI-PMH Harvester, for example, has been designed and developed as a MAS (Figure 7 depicts the Manent MAS architecture for the harvester component). The upper part of the architecture depicts all the agents designed for the Manent harvester. An AgentController is in charge of coordinating the agents towards the execution of a harvesting workflow. The AgentController knows at each moment the state of each agent and delivers communications that agents want to exchange between them. In particular, a ProxyAgent mediates all the communications between the AgentController and the other agents. Each agent type in the harvester is devoted to a specific task (behavior) that is carried out in parallel with other tasks (behaviors) executed by other agent types. The bottom part of the architecture depicts the behaviours of agents, which represent parallel tasks that strictly adhere to the OAI-PMH protocol directives.

The Manent service that performs text classification will be based on the Indiana Ontology concepts. It will implement a function f that takes, besides the Indiana Ontology concepts, a text as input and produces as output a list of pairs of the form $\langle$*concept*, *real value*$\rangle$, such that

$$f(t, o) = \{\textit{all pairs } \langle c \in o, r \in \mathbb{R} \rangle \textit{ such that } t \textit{ deals with } c \textit{ with confidence } r\}$$

Preliminary experiments have been conducted with Manent on a set of metadata records coming from different repositories. Such records were classified with a general purpose ontology based on the WordNet Domains[6] classification scheme. Results are reported in [37] and [38]. The same approach will be extended to all textual material in Manent, based on the Indiana Ontology. The classification will be straightforwardly applied and stored for different language versions of the same text, by exploiting the relationship between the English version of such text and its French or Italian versions. This information will be provided, if it exists, for every digital object inside Manent, included the ones in the Indiana GioNS DL that have been described as specific language versions of the same textual digital object.

A multilanguage retrieval service will then be provided by Manent, for each digital library that has been harvested, against a user query that wants to search text material related to but besides the one inside the Indiana GioNS DL. In this scenario, the user will ask for digital objects of other digital libraries that have been tagged with the Indiana Ontology concepts, i.e., that are pertinent to the Indiana MAS domain. A list of digital objects pertinent to the search criteria will then be returned.

In addition, the Manent interface will be accessible by the users through the Indiana MAS main GUI, and a list of available digital libraries will be provided to the user. For each of them, the user will be able to browse their structure, their metadata, and each digital object.

Manent has been designed as a Digital Library Management System, with an infrastructure able to rely on big data management that run upon distributed frameworks such as the Hadoop file system [52] and Apache HBase[7]. By relying on such technologies, Manent is able to offer all the services described above in an effective and efficient way.

4. **Interpretation as Cooperation among Agents and Holons: a Concrete Example.** In this section we exemplify the behavior of Indiana MAS by discussing how a sketch or an image can be correctly interpreted by AgentSketch and ImageRec, exploiting information previously discovered by other agents in the MAS.

Let us suppose that the user inputs either the image shown in Figure 8 (Mount Bego's engravings ZIV.GII.R19C.no 12 and 13 [17]) or the sketch shown in Figure 9 to the system.

[6] `http://wndomains.fbk.eu/`.

[7] `http://hbase.apache.org/`

FIGURE 8. Tracing of Mount Bego's engravings ZIV.GII.R19C.no 12 and 13.

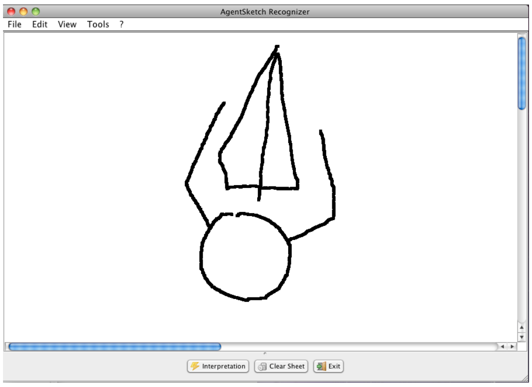

FIGURE 9. Sketch of Mount Bego's engravings ZIV.GII.R19C.no 12 and 13.

AgentSketch and ImageRec are both able to correctly recognize the corniculate at the bottom with a high confidence, but recognition of the element between the corniculate's horns fails: the shape in the image/sketch does not meet the standard dagger shape, and the confidence that could be a dagger is below the usually acceptable threshold. On the other hand, the confidence that it is a sinuous line, a canal watering the earth, a human, or any other shape that can be found in Mount Bego's engravings, as computed by the Symbol Recognizers in charge for these kinds of shapes, is even lower than the confidence that it is a dagger.

The process that leads to a correct interpretation despite this failure is the same for both AgentSketch and ImageRec, so we use AgentSketch to continue our example.

When sketches are as complex as rock carvings, the context to take into account to provide a correct and precise interpretation cannot be any longer limited to neighboring strokes, and cannot be hard-wired within the AgentSketch Symbol Interpretation Agent.

To make an example, in [17] a systematic analysis of meaningful associations between symbols characterizing Mount Bego's engravings has been carried out, leading to the identification of recurrent patterns such as:

- the dagger, symbol of the light, placed over a reticulate symbolizing the earth;
- the dagger above a corniculate, between its horns (interpretation: the god of storms that fertilizes the earth by means of the rain);
- sinuous lines symbolizing the stream;
- a water canal close to the horns of a corniculate (interpretation: water fertilizing the earth);
- a bull below a feminine figure with open arms and legs (interpretation: the myth known as the "son-husband", a symbolic representation of the high goddess or mother goddess giving birth to the bull, who in turn fertilizes the goddess).

Hard-wiring all such patterns within AgentSketch in order to reinforce, for example, the hypothesis that a given shape is a bull if it is below a feminine shape, is not feasible: relationships among symbols are too many, too complex and too dynamic, due to the many co-existing hypotheses made by archaeologists, and to the new ones that are continuously proposed.

The solution implemented in Indiana MAS decouples the discovery and representation of such patterns among symbols from their exploitation during the sketch and image interpretation stages: the system operates off-line to mine patterns like the ones described above from all the Indiana MAS sources.

Discovered patterns allow us to easily define and record new semantic information.

AgentSketch can use such information during its interpretation work: for every sketch, AgentSketch accesses the stored patterns involving symbols that have been recognized so far, whenever it needs to disambiguate the interpretation of a sketch. Patterns that support a given interpretation, make that interpretation stronger. The strongest interpretation is proposed to the Query Manager.

Let us go back now to the sketch shown in Figure 9. AgentSketch looks for patterns involving corniculates, and finds the following ones:

- *Stream_or_canal_watering_the_earth* **closeTo** *Corniculates*
- *Daggers* **above** *Corniculates*
- *Corniculates* **closeTo** *SinuousLines*

These patterns make the interpretation of the unknown shape as a Stream or a Canal, a Dagger, or a Sinuous Line feasible, but since *confidence(Dagger)* > *confidence(Stream_or_canal_watering_the_earth)* and *confidence(Dagger)* > *confidence(SinuousLine)*, the interpretation is that the sketch represents a dagger between the horns of a corniculate.

A system able to relate information coming from sources as heterogeneous as images, sketches, and texts in multiple languages, and to suggest interpretations to sketches and images based on such a large corpus of digital objects, may help archaeologists spread all over the world in finding useful relationships among the data they own and those made available by other researchers. For example, since the most common way to date engravings is to compare the engraving's subject to a physical artifacts whose date is known, Indiana MAS would help archaeologists in comparing images of physical artifacts and images of engravings, to make their dating more precise. Also, since similar artifacts can be found in very distant places, discovering strong similarities between artifacts (be them engravings or physical objects) by comparing and interpreting their images could open to new fascinating hypotheses on commercial and cultural contacts among ancient populations.

5. **Related and Future Work.** Several projects are funded by the European Commission with the aim of improving the quality and effectiveness of ICT in the cultural heritage field. They range from the adoption of well defined standards for resource description (the CIDOC CRM available at `http://www.cidoc-crm.org/`, and Europeana Data Model available at `http://pro.europeana.eu/`) to the definition of reference models for digital libraries (the FP6-IST *BRICKS* project available at `http://www.brickscommunity.org/`, and the *DELOS* network of excellence available at `http://www.delos.info/`). Other projects focus on the provision of multi-modal and multilingual services (e.g. the *MultiMatch* project available at `http://www.multimatch.eu/`). In the rock art field, the RAMP, Rock Art Mobile Project [49], is a social and mobile application designed and developed for the interpretation of Rock Art in different UK areas. The project focuses on the aspects of interpretation seen as a seamless process of interaction between users and designers. Starting from the provision of a on-line archive of rock art cards accessible on

mobile devices, an innovative communication space is built where to collect visiting rock art impressions, hence improving the design of tools that support the rock art experience.

An awareness campaign for promoting rock art experience of UNESCO protected world heritage site of Val Camonica in Northern Italy was the purpose of the project reported in [51]. A collaborative computer game based on virtual engravings for multi-touch display devices was the outcome of the project.

The digital archive available at `http://www.digitalrockart.org/` collects thousands of digital images and forms from several rock arts sites, stored on a volunteer basis. The goal of the tool is to promote the recording and storage of such sites and make them available on-line for the development of applications based on such recordings, for both researchers and users. Information is recorded in different formats grouped in several views: from narratives to pictures, from sites maps to click-able paths of hyperlinks, and so on. Other rock art galleries are preserved in Oregon, USA (available at `http://www.oregonrockart.com/gallery.htm`), and in the RAMP project website (the Stan Beckensall archive, available at `http://rockart.ncl.ac.uk/`).

The digital preservation, classification, and interpretation of rock carvings raises many scientific challenges, such as the integration of data coming from multiple sources, and the interpretation of drawings whose meaning may vary based on contextual information. The Indiana MAS project tackles all such issues by exploiting intelligent agents that ensure the required degree of flexibility and autonomy in a highly dynamic and heterogeneous environment. In this sense, it is much more ambitious than other projects in the Rock Art domain, since it takes multimedia and multilingual objects into account and deserves a great attention to their storage and classification using state of the art technologies and standards, whereas most similar projects just focus on one specific type of objects, and neglect multilingual issues.

Currently we are integrating the developed components within Indiana MAS and, although, the fully implementation and testing of this integration is our main close-future work, other improvements to the Indiana MAS components are foreseen as well. In particular, we are working on making AgentSketch more modular and dynamic by allowing SRAs to read the specification of the symbols they have to recognize from the Indiana Ontology, instead of hard-wiring their definition into LADDER rules. Also, we will exploit the results of the Ontologica project carried out with Ansaldo STS, the Italian leader in railways monitoring and signaling [7], to make the user capable of interacting with the Query Agent using natural language queries. Finally, new algorithms for sketch, image and text understanding, and new proposals for supporting agents coordination and interaction, will be produced as results of the basic research funded by the Indiana MAS project.

Acknowledgments. This research is supported by the "Indiana MAS and the Digital Preservation of Rock Carvings" FIRB project funded by the Italian Ministry for Education, University and Research, under grant RBFR10PEIT.
The authors are grateful to Francesca Odone for her helpful comments on Section 3.2.2.

REFERENCES

[1] C. Alvarado and R. Davis. Sketchread: a multi-domain sketch recognition engine. In *Procs. of UIST'04*, pages 23–32. ACM, 2004.
[2] D. Avola, P. Bottoni, A. Dafinei, and A. Labella. Fcbd: An agent-based architecture to support sketch recognition interfaces. In *Procs. of DMS*, pages 295–300, 2011.
[3] S. Azar, L. Couvreury, V. Delfosse, B. Jaspartz, and C. Boulanger. An agent-based multimodal interface for sketch interpretation. In *Proceedings of International Workshop on Multimedia Signal Processing*, 2006.

[4] Z. A. Baig. Review: Multi-agent systems for protecting critical infrastructures: A survey. *J. Netw. Comput. Appl.*, 35(3):1151–1161, May 2012.

[5] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(4):509–522, 2002.

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

[7] D. Briola, R. Caccia, M. Bozzano, and A. Locoro. Ontologica: Exploiting ontologies and natural language for representing and querying railway management logics. In M. Graña, C. Toro, J. Posada, R. J. Howlett, and L. C. Jain, editors, *KES*, volume 243 of *Frontiers in Artificial Intelligence and Applications*, pages 376–385. IOS Press, 2012.

[8] L. Candela, G. Athanasopoulos, D. Castelli, K. El Raheb, P. Innocenti, Y. Ioannidis, A. Katifori, A. Nika, G. Vullo, and S. Ross. The digital library reference model. Technical Report Deliverable D3.2b, DL.org: Coordination Action on Digital Library Interoperability, Best Practices and Modelling Foundations. Funded under the 7th Framework Programme. Project Number 231551., 2011.

[9] G. Casella, V. Deufemia, V. Mascardi, G. Costagliola, and M. Martelli. An agent-based framework for sketched symbol interpretation. *J. Vis. Lang. Comput.*, 19(2):225–257, 2008.

[10] G. Casella, V. Deufemia, V. Mascardi, M. Martelli, and G. Tortora. Reasoning about hand-drawn sketches: An approach based on intelligent software agents. In M. Sebillo, G. Vitiello, and G. Schaefer, editors, *VISUAL 2008*, volume 5188 of *Lecture Notes in Computer Science*, pages 302–314. Springer, 2008.

[11] L. Chen and K. P. Sycara. Webmate: A personal agent for browsing and searching. In *Agents*, pages 132–139, 1998.

[12] D. Choinski and M. Senik. Ontology based knowledge management and learning in multi-agent system. In *Agent and Multi-Agent Systems. Technologies and Applications*, Lecture Notes in Computer Science, pages 65–74. Springer, 2012.

[13] M. Cossentino, N. Gaud, V. Hilaire, S. Galland, and A. Koukam. ASPECS: an agent-oriented software process for engineering complex systems. *Autonomous Agents and Multi-Agent Systems*, 20(2):260–304, 2010.

[14] G. Costagliola, V. Deufemia, and M. Risi. Using error recovery techniques to improve sketch recognition accuracy. In W. Liu, J. Lladós, and J.-M. Ogier, editors, *Graphics Recognition. Recent Advances and New Opportunities*, pages 157–168. Springer-Verlag, 2008.

[15] M. Şensoy and P. Yolum. Automating user reviews using ontologies: an agent-based approach. *World Wide Web*, 15(3):285–323, May 2012.

[16] DCMI. Dublin Core Metadata Initiative.

[17] H. de Lumley and A. Echassoux. 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, 2009.

[18] H. V. de Sompel, M. Nelson, C. Lagoze, and S. Warner. Resource harvesting within the OAI-PMH framework. *D-Lib Magazine*, 10(12), 2004.

[19] S. M. Deen, editor. *Agent Based Manufacturing – Advances in the Holonic Approach.* Advanced Information Processing. Springer, 2003.

[20] D. Derbyshire, I. Ferguson, J. Müller, M. Pischel, and M. Wooldridge. Agent-based digital libraries: Driving the information economy. In *Procs of the 6th Workshop on Enabling Technologies on Infrastructure for Collaborative Enterprises*, WET-ICE '97, pages 82–86. IEEE Computer Society, 1997.

[21] V. Deufemia, L. Paolino, and H. de Lumley. Petroglyph recognition using self-organizing maps and fuzzy visual language parsing. In *Proceedings of the IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2012)*, pages 852–859, 2012.

[22] V. Deufemia, L. Paolino, G. Tortora, A. Traverso, V. Mascardi, M. Ancona, M. Martelli, N. Bianchi, and H. de Lumley. Investigative analysis across documents and drawings: visual analytics for archaeologists. In *Proceedings of the International Working Conference on Advanced Visual Interfaces*, pages 539–546. ACM, 2012.

[23] D. Fernández-Pacheco, J. Conesa, and N. Aleixos. A new agent-based paradigm for recognition of free-hand sketches. *Procedia Computer Science*, 1(1):2013 – 2022, 2010.

[24] K. Fischer. Agent-based design of holonic manufacturing systems. *Journal of Robotics and Autonomous Systems*, 27(1-2):3–13, 1999.

[25] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. In P. M. B. Vitányi, editor, *EuroCOLT'95*, volume 904 of *Lecture Notes in Computer Science*, pages 23–37. Springer, 1995.

[26] E. Friedman-Hill. Jess, the rule engine for the Java platform. Sandia National Laboratories, Livermore, CA. Online at `http://www.jessrules.com/`. Last accessed: February the 15th, 2012.

[27] C. Gerber, J. H. Siekmann, and G. Vierke. Holonic multi-agent systems. Technical Report DFKI-RR-99-03, Deutsches Forschungszentrum für Künztliche Inteligenz - GmbH, Postfach 20 80, 67608 Kaiserslautern, FRG, 1999.

[28] T. R. Gruber. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5:199–220, 1993.

[29] M. Hadzic, P. Wongthongtham, T. S. Dillon, and E. Chang. *Ontology-Based Multi-Agent Systems*, volume 219 of *Studies in Computational Intelligence*. Springer, 2009.

[30] T. Hammond and R. Davis. LADDER, a sketching language for user interface developers. *Computers & Graphics*, 29(4):518–532, 2005.

[31] T. Hammond and R. Davis. Tahuti: a geometrical sketch recognition system for UML class diagrams. In *ACM SIGGRAPH 2006 Courses*. ACM, 2006.

[32] B. Horling and V. Lesser. A survey of Multi-Agent organizational paradigms. *The Knowledge Engineering Review*, 19(4):281–316, 2005.

[33] F.-S. Hsieh. Collaborative workflow management in holonic multi-agent systems. In J. O'Shea, N. T. Nguyen, K. A. Crockett, R. J. Howlett, and L. C. Jain, editors, *Agent and Multi-Agent Systems. Technologies and Applications*, volume 6682 of *Lecture Notes in Computer Science*, pages 383–393. Springer, 2011.

[34] R. Juchmes, P. Leclercq, and S. Azar. A freehand-sketch environment for architectural design supported by a multi-agent system. *Computers & Graphics*, 29(6):905 – 915, 2005.

[35] T. H. W. Lam, J. N. K. Liu, and R. S. T. Lee. Master: an intelligent ontology-based multi-agent system for sightseer. *International Journal of Software Engineering and Knowledge Engineering*, 19(2):137–157, 2009.

[36] H. Lieberman. Letizia: An agent that assists web browsing. In *14th Int. Joint Conference on Artificial Intelligence*, pages 924–929. Morgan–Kaufmann, 1995.

[37] A. Locoro, D. Grignani, and V. Mascardi. MANENT: An infrastructure for integrating, structuring and searching digital libraries. In M. Biba and F. Xhafa, editors, *Learning Structure and Schemas from Documents*, volume 375 of *Studies in Computational Intelligence*, pages 315–341. Springer, 2011.

[38] A. Locoro, D. Grignani, and V. Mascardi. When you doubt, abstain: From misclassification to epoché in automatic text categorisation. In J. F. Hübner, J.-M. Petit, and E. Suzuki, editors, *Web Intelligence/IAT Workshops*, pages 209–212. IEEE Computer Society, 2011.

[39] A. Locoro, V. Mascardi, D. Briola, M. Martelli, M. Ancona, V. Deufemia, L. Paolino, G. Tortora, G. Polese, and R. Francese. The Indiana MAS project: Goals and preliminary results. In *Proceedings of the 13th Workshop on Objects and Agents*, volume 892 of *CEUR Workshop Proceedings*. CEUR-WS.org, 2012.

[40] D. G. Lowe. Object recognition from local scale-invariant features. In *International Conference on Computer Vision*, pages 1150–1157. IEEE Computer Society, 1999.

[41] G. Mackenzie and N. Alechina. Classifying sketches of animals using an agent-based system. In N. Petkov and M. Westenberg, editors, *Computer Analysis of Images and Patterns*, volume 2756 of *Lecture Notes in Computer Science*, pages 521–529. Springer Berlin Heidelberg, 2003.

[42] V. Mascardi, V. Deufemia, D. Malafronte, A. Ricciarelli, N. Bianchi, and H. de Lumley. Rock art interpretation within Indiana MAS. In G. Jezic, M. Kusek, N.-T. Nguyen, R. Howlett, and L. Jain, editors, *Agent and Multi-Agent Systems. Technologies and Applications*, volume 7327 of *Lecture Notes in Computer Science*, pages 271–281. Springer, 2012.

[43] V. Mascardi, J. Hendler, and L. Papaleo. Semantic web and declarative agent languages and technologies: Current and future trends. In *DALT X: Revised, Selected and Invited Papers*, Lecture Notes in Computer Science, pages 197–202. Springer, 2012.

[44] T. Ouyang and R. Davis. Learning from neighboring strokes: Combining appearance and context for multi-domain sketch recognition. In *Advances in Neural Information Processing Systems*, pages 1401–1409, 2009.

[45] A. Rodríguez-González, J. Torres-Niño, G. Hernández-Chan, E. Jiménez-Domingo, and J. M. Alvarez-Rodríguez. Using agents to parallelize a medical reasoning system based on ontologies and description logics as an application case. *Expert Syst. Appl.*, 39(18):13085–13092, Dec. 2012.

[46] B. Rubin. Radon, cosine and sine transforms on real hyperbolic space. *Advances in Mathematics*, 170(2):206–223, 2002.

[47] S. Salzberg. On comparing classifiers: Pitfalls to avoid and a recommended approach. *Data Min. Knowl. Discov.*, 1(3):317–328, 1997.

[48] J. Sánchez, J. Leggett, and J. Schnase. Ags: introducing agents as services provided by digital libraries. In *Procs. of DL '97*, pages 75–82, 1997.

[49] J. Sánchez, J. Leggett, and J. Schnase. Situating cultural technologies outdoors: Designing for mobile interpretation of rock art in rural Britain. In *Procs of the annual conference on Museums and the Web*, 2011.

[50] M. Seidl and C. Breiteneder. Detection and classification of petroglyphs in gigapixel images – preliminary results. In *VAST: International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage - Short and Project Papers*, pages 45–48, 2011.

[51] M. Seidl, P. Judmaier, F. Baker, C. Chippindale, U. Egger, N. Jax, C. Weis, M. Grubinger, and G. Seidl. Multi-touch rocks: Playing with tangible virtual heritage in the museum – first user tests. In *VAST: International Symposium on Virtual Reality, Archaeology and Intelligent Cultural Heritage - Short and Project Papers*, pages 73–76, 2011.

[52] K. Shvachko, H. Huang, S. Radia, and R. Chansler. The Hadoop distributed file system. In *Procs of the 26th IEEE Symposium on Massive Storage Systems and Technologies (MSST2010)*. IEEE Computer Society, 2001.

[53] M. Smith, M. R. Barton, M. Branschofsky, G. McClellan, J. H. Walker, M. J. Bass, D. Stuve, and R. Tansley. DSpace: An open source dynamic digital repository. *D-Lib Magazine*, 9(1), 2003.

[54] P. A. Viola and M. J. Jones. Rapid object detection using a boosted cascade of simple features. In *CVPR 2001*, pages 511–518. IEEE Computer Society, 2001.

[55] S.-Y. Yang, C.-L. Hsu, D.-L. Lee, and L. Y. Deng. Faq-master: an ontological multi-agent system for web faq services. *WSEAS Trans. Info. Sci. and App.*, 5(3):221–228, Mar. 2008.

[56] Y. Yang, O. Rana, D. Walker, C. Georgousopoulos, G. Aloisio, and R. Williams. Agent based data management in digital libraries. *Parallel Comput.*, 28(5):773–792, May 2002.

[57] S. Zambanini and M. Kampel. Automatic coin classification by image matching. In F. Niccolucci, M. Dellepiane, S. P. Serna, H. E. Rushmeier, and L. J. V. Gool, editors, *VAST 2011*, pages 65–72. Eurographics Association, 2011.

[58] X. Zhang and D. Norrie. Holonic control at the production and controller levels. In *2nd Int'l Workshop Intelligent Manufacturing Systems, IMS'99*, pages 215–224, 1999.

[59] Q. Zhu, X. Wang, E.Keogh, and S. Lee. An efficient and effective similarity measure to enable data mining of petroglyphs. *Data Mining and Knowledge Discovery*, 23:91–127, 2011.