

*OntoScene, A Logic-Based Scene Interpreter: Implementation and Application in the Rock Art Domain**

DANIELA BRIOLA

*Department of Computer Sciences, Systems and Communications
University of Milano Bicocca, Italy
(e-mail: daniela.briola@unimib.it)*

VIVIANA MASCARDI and MASSIMILIANO GIOSEFFI

*Department of Informatics, Bioengineering, Robotics, and Systems Engineering
University of Genova, Italy
(e-mails: viviana.mascardi@unige.it, gmaxsun89@gmail.com)*

submitted 14 January 2019; revised 5 November 2019; accepted 12 December 2019

Abstract

We present OntoScene, a framework aimed at understanding the semantics of visual scenes starting from the semantics of their elements and the spatial relations holding between them. OntoScene exploits ontologies for representing knowledge and Prolog for specifying the interpretation rules that domain experts may adopt, and for implementing the SceneInterpreter engine. Ontologies allow the designer to formalize the domain in a reusable way and make the system modular and interoperable with existing multiagent systems, while Prolog provides a solid basis to define complex rules of interpretation in a way that can be affordable even for people with no background in Computational Logics. The domain selected for experimenting OntoScene is that of prehistoric rock art, which provides us with a fascinating and challenging testbed.

KEYWORDS: Prolog, ontologies, multiagent systems, visual languages, scene interpretation

1 Introduction

Human perception of complex visual scenes has been studied for a long time in psychology and neuroscience ([Kondo *et al.* 2017](#)): according to the seminal work on “high-level scene perception” ([Henderson and Hollingworth 1999](#)), besides low-level or early vision, concerned with extraction of physical properties such as depth, color, and texture from an image ([Marr 1982](#)), and intermediate-level vision, concerned with extraction of shape

* We thank Prof. Henry de Lumley and Annie Echassoux for granting us the permission to reproduce some figures from their book ([de Lumley and Echassoux 2011](#)), and Martine Bertéa, Rights Director of CNRS éditions, for helping us in obtaining their permission. We are grateful to Dr. Nicoletta Bianchi for her precious support in the IndianaMAS project and in the activities we faced after its conclusion. Finally, we thank the anonymous reviewers for their thorough reading and for their constructive comments.

and spatial relations that can be determined without regard to meaning (Ullman 1996), a further level of vision is required to perceive and understand a scene:

high-level vision concerns the mapping from visual representations to meaning and includes [...] the identification of objects and scenes.

In their recent studies, Kveraga and Bar (2014) and Baldassano (2015) demonstrate that the brain has regions related to higher-order properties like overall geometry, interactions between objects, esthetic beauty, or memorability of a scene. These regions show a larger response to full scenes than to isolated objects.

Artificial intelligence can play a major role in modeling and understanding, on the one hand, and reproducing, on the other, the way visual scenes are interpreted by humans.

While deep learning has shown impressive potential in recognizing images (He *et al.* 2016; Simonyan and Zisserman 2014; Wan *et al.* 2014; Donahue *et al.* 2014), hence providing an ideal tool for low-level and intermediate-level vision, tackling the high-level vision, and associating a meaning with complex scenes may require an explicit and symbolic representation of the domain knowledge, and the ability to reason over it.

To understand the semantics of a scene starting from the semantics of its elements and the relations holding among them we developed OntoScene, which exploits a powerful combination of ontologies and Prolog: ontologies are used for representing knowledge, and Prolog for specifying the rules that domain experts actually use to interpret visual scenes and for implementing the SceneInterpreter engine. OntoScene also relies upon technologies developed in the multiagent systems (MASs) area: it is in fact part of a holonic MAS (Gerber *et al.* 1999) named IndianaMAS (Mascardi *et al.* 2014; Briola *et al.* 2014; Briola 2016; Briola *et al.* 2017) where agents and MASs devoted to multilingual text understanding, hand-drawn sketch recognition, human interaction, and integration of digital libraries, cooperate and coordinate with the OntoScene framework to classify heterogeneous digital objects.

Following a widely accepted approach for the interpretation of a scene, we consider a scene as an instance or phrase of a visual language where, by analogy with textual languages, relevant graphical symbols can be understood as lexical components or tokens that can be aggregated through the syntactic rules defined according to relations holding among them. Tokens are the sub-images that make up the scene, the grammar is represented by rules defined by the domain expert, and geometric relationships are “vertical”, “overlapping”, “close”, and the like, and represent aggregation operators. To allow domain experts to describe the rules for interpreting scenes using a language close to the one in which these rules would be expressed in natural language, we use Prolog. We have designed a user-friendly language that domain experts may use. This rule-based, domain-specific language is very similar to Prolog but it hides most Prolog technicalities and can be compiled into standard Prolog clauses.

OntoScene consists of:

- Detector and Classifier, two external modules (whose functioning is outside the scope of this paper, and which could be based on our own previous proposals (Briola *et al.* 2017) or on more recent deep learning techniques) that partition the input image into tokens and associate a list of classifications with them, respectively;
- SceneInterpreter, the Prolog core of OntoScene; it reasons on a symbolic representation of images that make up a scene and returns their interpretations;

- OntoScene Agent, an agent providing the interface between OntoScene and the other agents in IndianaMAS;
- The OntoScene Ontology, which models general concepts needed by OntoScene to work, as well as domain-dependent concepts.

To show the potentiality of the OntoScene framework and to verify the concrete applicability of the proposed solution, we exploit it for the interpretation of complex scenes from the rock art domain, in particular the one of Mount Bego, in France: Mount Bego archeological site is well known for its petroglyphs (carvings on rocks), ancient testimonies of human first activities (Bianchi 2011; Bicknell 1913; de Lumley and Echassoux 2009; 2011). These carvings represent animals, geometric shapes, rural elements, and anthropomorphic figures, often represented together to form complex scenes: if identifying and interpreting single elements could be quite simple, interpreting complex scenes requires a very detailed knowledge of the domain and offers a challenging testbed to OntoScene.

The core functionality of SceneInterpreter, namely the generation of all the possible scene interpretations according to the interpretation rules, is implemented by Donald Knuth's *Algorithm X* for the exact cover problem (Knuth 2000). Algorithm X is a state space searching algorithm that natively exploits depth-first search and backtracking: Prolog turns out to be the perfect language for its implementation. Also, Prolog is very effective as a scene interpretation rule modeling language. Such rules are either sketched by the domain experts using the user-friendly syntax that we devised to mask Prolog details or written by ourselves in close cooperation with the experts: in both cases, the domain expert that we involved in the experiments, the archeologist Dr. Nicoletta Bianchi, easily grasped the concepts of unification and backtracking, that allowed her to specify the rules she had in mind, often based on a generate and test technique, in a natural and intuitive way.

The paper is organized as follows: Section 2 offers the background knowledge needed for reading the paper and overviews works related to ours; Section 3 provides a gentle introduction to OntoScene; Section 4 describes how we modeled domain and spatial knowledge; Section 5 presents the SceneInterpreter module and exemplifies its functioning on a synthetic domain; Section 6 describes the experiments carried out in the rock art domain; Section 7 concludes and outlines the future directions of our research.

2 Background and related work

OntoScene is used inside the IndianaMAS holonic MAS, which has been designed and developed as a JADE (Java Agent DEvelopment Framework (Bellifemine et al. 2007)) MAS. Although OntoScene's main components are not agents, its interface toward the IndianaMAS components is the JADE OntoScene Agent, which heavily exploits the tools that JADE offers to integrate ontologies in the MAS. Assuming the reader is familiar with knowledge representation in general and with ontologies in particular¹, and in Section 2.1, we provide a brief introduction to IndianaMAS, to JADE, and to the way ontologies are supported therein. We also provide references to the JPL library²

¹ The reader may find an introduction to computational ontologies in (Guarino et al. 2009), the specification of the OWL Web Ontology Language in (McGuinness et al. 2004), and information on Protégé on the official website, <https://protege.stanford.edu/>, accessed on July 2019.

² <http://www.swi-prolog.org/packages/jpl/>, accessed on July 2019.

for interfacing SWI-Prolog and Java, and to the JTS Topology Suite³ we used to compute relationships among elements in a scene. Section 2.2 compares our work with related proposals in the logic-based visual languages field, and with spatial ontologies and ontology-driven scene interpretation.

2.1 Background

IndianaMAS. According to the seminal paper by Wooldridge and Jennings (1995), an agent is a hardware or, more usually, a software-based computer system that is autonomous (agents operate without the direct intervention of humans or others and have some kind of control over their actions and internal state); social (agents interact with other agents and possibly humans via some kind of agent communication language); situated and reactive (agents perceive their environment and respond in a timely fashion to changes that occur in it); and proactive (agents do not simply act in response to their environment, but they are able to exhibit goal-directed behavior by taking the initiative).

Agents are the right tool for coordinating the functioning of software artifacts that show different capabilities and are possibly distributed across a network, with the purpose of making the software architecture as modular, flexible, and reliable as possible.

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 (“IndianaMAS” for short⁴), funded by the Italian Ministry for Education, University and Research (MIUR) and spanning from March 2012 to February 2015, developed 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. IndianaMAS enables the preservation of all kinds of available data about rock carvings, such as images, geographical objects, textual descriptions of the represented subjects, allowing the domain experts to organize and structure such digital objects in a standard way and to supply domain experts with facilities for issuing complex queries on the data repositories.

The choice of agent technology for addressing the IndianaMAS goals was a very natural one, given the need that each component of the system, while operating in a highly autonomous way, could interact and coordinate with the other components to share information and to reason about it in the most effective way. As discussed by Mascardi *et al.* (2014), the three key services offered by IndianaMAS (sketch recognition, image recognition, and multilingual access to digital libraries) are provided by systems that may be MASs themselves and that are seen as black boxes by the IndianaMAS agents.

Besides OntoScene, the main components of IndianaMAS, sketched in Figure 1, are:

- The Indiana Ontology, which structures the domain of interest, consists of sub-ontologies, among which the OntoScene Ontology⁵ and is accessed by all the agents and components in the system.

³ <https://locationtech.github.io/jts/>, accessed on July 2019.

⁴ We use “IndianaMAS” to denote both the funded project and the MAS that resulted from it. The project website, <http://indianamas.disi.unige.it/>, accessed on July 2019, gives access to all the project’s deliverables and papers.

⁵ The OntoScene Ontology shares the “Thing” class with the Indiana Ontology: we will not address the question of whether it should be named an “ontology” or a “sub-ontology”, as this distinction is not relevant for the paper, and we will always use “ontology” to describe it.

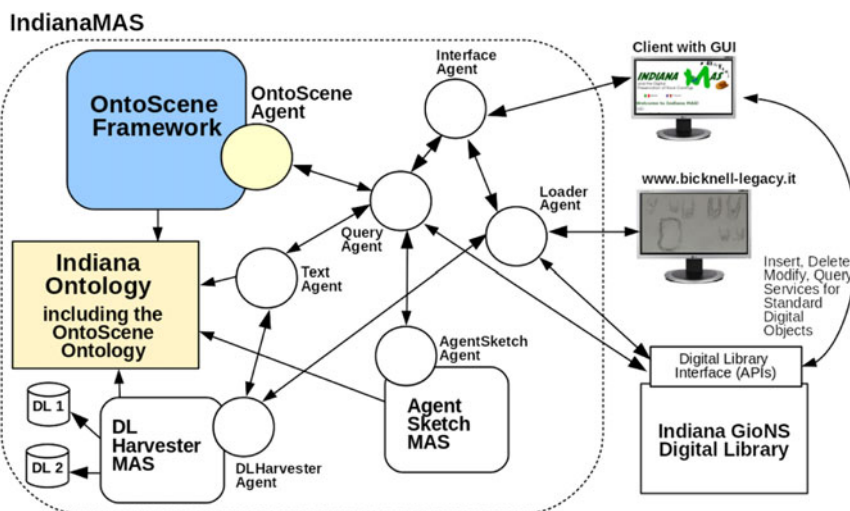


Fig. 1. IndianaMAS architecture and data flow.

- Client with a graphical user interface, for interacting with IndianaMAS.
- The Indiana GioNS Digital Library, which contains all the digital objects inserted into the system by registered users, together with their metadata, needed for their later retrieval.
- Text Agent able to interpret multilingual documents according to the Indiana Ontology.
- Query Agents, each managing one query coming from the client.
- Loader Agent collecting new data from external resources like the Bicknell Legacy website⁶ and managing the creation and insertion of new digital objects into the Indiana GioNS Digital Library.
- Interface Agent, managing the creation of new Query Agents.
- The Digital Library (DL) Harvester MAS, which independently and proactively searches digital libraries on the web to retrieve new images and texts related to the domain modeled by the Indiana Ontology.
- AgentSketch MAS, which interprets manual drawings based on the Indiana Ontology.

JADE. JADE is a Java-based software platform that supports the development of agents and MASs thanks to a graphical user interface and tools supporting the MAS debugging and deployment phases. JADE MASs can be distributed across machines in a way that is fully transparent to the developer. The minimal system requirement is the Java runtime environment or JDK, version 5.

Ontologies in JADE. JADE helps developers in achieving semantic interoperability between agents thanks to a simple and fast way to exploit ontologies directly inside the platform and the agents: agents can exchange messages referring to a shared ontology,

⁶ <http://www.bicknell-legacy.it>, accessed on July 2019.

and then rely on the JADE Ontology management offered by the ContentManager class. The developer may use an ontology to formalize what the agents know (Concepts and Predicates) and can do (Actions), and share this ontology among the agents: in this way, knowledge is modeled outside the agents, boosting modularity and reuse, and the content of messages is based on a shared ontology, facilitating interactions and simplifying the serialization phase that is then demanded to the JADE platform.

The three types of objects considered when creating an ontology for JADE are:

- Predicates: boolean expressions describing something about the agent environment or its beliefs.
- Concepts: structured objects describing the elements of the world and their relationships.
- Agent actions: special Concepts modeling what an agent can do and can be requested to do with a message.

If two agents share the same ontology, one agent can request the other to perform an Action and can receive an answer containing the Action results, which will be a Concept, a Predicate, or a list of them.

To allow developers to automatically generate a Java representation of an OWL ontology coherent with the JADE requirements, the tool OntologyBeanGenerator⁷ can be adopted. The latest version of OntologyBeanGenerator available on the official website is 4.1, including a basic ontology modeling the over mentioned concepts: domain-specific concepts must be added as subclasses of Concept, Predicate, and Agent, and then the tool will provide a Java representation of the ontology, directly usable by JADE.

Given some limitations of that version, we developed OntologyBeanGenerator 5.0 (Briola *et al.* 2018) as a new Protégé plugin⁸. OntologyBeanGenerator 5.0 (OBG5.0 in the sequel, available from www.disi.unige.it/person/MascardiV/Download/OBG5.0.zip) has been developed with three goals in mind: correcting some bugs of OntologyBeanGenerator 4.1; adding the methods and exceptions management directly inside the ontology; and producing an additional output to support the OntoScene framework.

The main improvements of OBG5.0 w.r.t. OntologyBeanGenerator 4.1 are:

- Addition of a new tab called Java Method Mapper to manage the methods creation and exportation: the purpose of the new tab is to offer the designer of the ontology a way to directly add methods to the Java version of the ontology: in the previous version, the only option was to add a property and consequently to get the setter and getter automatically.
- Exception management: methods are allowed to raise Exceptions. To do this, a specific ontology to be imported has been created. Thanks to this addition, Exceptions can be exchanged between agents too, since they are a subclass of Concept.
- Correction of some bugs that were present in the ontology generation stage.
- Possibility to export the class hierarchy in a Prolog format: in order to implement Prolog rules that reason about the ontology, we need a Prolog representation

⁷ <https://protegewiki.stanford.edu/wiki/OntologyBeanGenerator>, accessed on July 2019.

⁸ We asked and obtained a written consent from the author of OntologyBeanGenerator, Chris van Aart, to extend the original source code.



Fig. 2. A simple ontology to be exported in Prolog.

of it. To achieve this goal, we added an automatic ontology export functionality to OBG5.0. The obtained Prolog representation only formalizes the classes hierarchy, as this is the only knowledge we currently need in OntoScene.

As an example, the Prolog version of the class hierarchy shown in Figure 2 is:

```

subclass_of('C1', 'C11').
subclass_of('C1', 'C12').
subclass_of('C11', 'C111').
subclass_of('C2', 'C21').
  
```

The JPL Library. JPL can be used to embed SWI-Prolog in Java as well as for embedding Java in SWI-Prolog. In both setups, it provides a bidirectional interface. The two predicates that we used for accessing Java from inside SceneInterpreter are:

`jpl_new(+X, +Params, -V)` where `X` is an object (non-array) type or descriptor and `Params` is a list of values or references unifies `V` with the result of an invocation of that type's most specifically typed constructor to whose respective formal parameters the actual `Params` are assignable (and assigned).

`jpl_call(+V, +Method, +Params, -Result)` unifies `Result` with a JPL reference to (or value of) the result of calling the named `Method` of `V` with `Params`.

The JTS suite. The JTS Topology Suite (JTS) is an open source Java software library that provides an object model for planar geometry together with a set of fundamental geometric functions. In OntoScene, it was used to implement the basic relations between regions that characterize the Region Connection Calculus (RCC, (Li and Ying 2003; Randell and Cohn 1989)) such as `disjoint`, named “Disconnected” in RCC, `overlap` (“Partially Overlapping” in RCC), and `contains` (“Non-Tangential Proper Part Inverse” in RCC), plus further derived relations.

Using JPL and JTS together. If `GR` is a reference to the implementation of the interface for geometric relations, `jpl_call(GR, contains, [BB1, BB2], @(true))` succeeds if the method `contains`, implemented in Java by exploiting the API offered by JTS, and called on `BB1` (where `BB` stands for the “Bounding Box” of one image) and `BB2` (the bounding box of another image), returns `true`. In a similar way, we may have the following predicate calls in a Prolog piece of code

- `jpl_call(GR, overlap, [BB1, BB2], @(true))`,
- `jpl_call(GR, vertical, [BB2, BB1, 'up'], @(true))`,
- `jpl_call(GR, near, [BB1, BB2, 0.5], @(true))` (0.5 is the threshold for considering two bounding boxes close to each other, expressed in pixels),
- ...

with their intuitive meaning, better explained in Section 4.2. As a more complex example, the call to `jpl_call(GR, group, [JavaBBs, 0.5], @(true))` works if `JavaBBs` is the Java representation of a Prolog list, and the `group` method called on that list with `0.5` as proximity threshold returns `true`.

2.2 Related work

Logic-based visual languages. Many approaches for dealing with visual languages have been proposed in the literature: this research area has a long tradition, with both an *ad hoc* conference established in 1984, VL/HCC⁹, and a high-quality journal, the Journal of Visual Languages and Computing.¹⁰ In this section, we review some approaches that use logical or relational formalisms for recognizing and understanding visual languages, starting from the older and more established ones, and moving toward more recent proposals. A complementary approach, which is out of the scope of this paper, is to use visual programming approaches to specify logic-based languages, as done by Ladret and Rueher (1991) and Agustí *et al.* (1998)

Defining visual languages using a logic-based language in general, and Prolog in particular, ensures that declarative and operational semantics can be shared among humans and between humans and machines. The declarative semantics allows both humans and machines to reason about the specification independently of the implementation, while the operational semantics allows the generation and recognition of images defined by the specification. After a very active period in the early nineties of the last century, the “logic-based visual languages” research field has produced less results, probably due to the raise of statistical approaches in the meanwhile.

Crimi *et al.* (1991) introduced the concept of relational grammars: while textual languages use an implicit sequential concatenation relationship, the proposed extension relaxes this constraint by providing an arbitrary number of geometric relationships. Helm and Marriott (1991) defined the relationships between images and their meaning via a class of declarative and constraint-based specification languages, written in Prolog, and Wittenburg *et al.* (1991) presented a formalism called *unification-based grammar* and a parsing algorithm for visual languages. The formalism extends D-PATR (Karttunen 1986) with logical constraints and a new bottom-up parsing method. Meyer (1992) introduced a new technique to extend logic programming with terms representing partially specified images. To this aim, the picture clause grammar, a form of specification for visual languages similar to the definite clause grammar of textual languages, is defined. None of these proposals come with an implemented prototype, making their practical applicability limited.

Santosh *et al.* (2009) proposal is close to ours both in the system architecture and in the methodological approach, but not in the final goal. They aim at expressing graphic symbols by a number of graphical primitives that may be of any complexity and connecting relationships that can be deduced from state-of-the art image treatment and analysis tools. The existence of suitable tools for image pre-processing is also assumed by us, by

⁹ The most recent edition of VL/HCC dates back to 2018, <https://vlhcc18.github.io/index.html>, accessed on July 2019.

¹⁰ <https://www.journals.elsevier.com/journal-of-visual-languages-and-computing>, accessed on July 2019.

including the Detector and Classifier modules presented in the next sections in the OntoScene architecture. The symbolic representation obtained by the image analysis tools is then provided to an inductive logic programming solver that outputs a set of logical rules that define the positive example set. On the contrary, we provide the symbolic representation of elements detected in the scene to a Prolog program that, thanks to rules that model the domain expert knowledge, provide a semantic interpretation of the scene.

Antanas *et al.* (2012) present a framework combining compositional hierarchies, qualitative spatial relations, relational instance-based learning, and robust feature extraction. For each layer in the hierarchy, substructures in the images are detected, classified, and then employed one layer up the hierarchy to obtain higher-level semantic structures, by making use of qualitative spatial relations implemented in Prolog. Given that we may have scenes that include scenes, we support a hierarchical structure as well. So far, we only employed two levels in the hierarchy (one scene that includes another scene, that only includes “atomic” tokens, as in Table 13, third and fourth images) but there are in principle no reasons for adding more layers. W.r.t. that work, we also have a domain ontology and a MAS coordinating the interactions among the framework components.

In their work, Di Martino and Esposito (2016) do not consider any low-level image processing stage, but integrate a domain ontology in the system architecture, like in OntoScene: the authors describe a procedure and a prototype implementation for the automatic recognition of design patterns from documentation of software artifacts design and implementation, provided in XMI¹¹. The procedure exploits a semantic representation of the patterns to be recognized, based on an existing ontology. Both the UML set of diagrams related to the analyzed software artifacts and the patterns represented in the ontology are translated into a Prolog knowledge base. A Prolog program implements the heuristics and features that trigger the recognition on that knowledge base.

Although not based on logic programming, it is worth mentioning the work by Hammond and Davis (2007), which uses the rule-based language Jess (Hill 2003) for specifying how sketched diagrams in a domain are drawn, displayed, and edited, and the work by Costagliola *et al.* (2005), which uses rules named “sketch patterns” for describing and recognizing diagrammatic sketch languages, and that are very close to Jess rules.

Spatial ontologies and ontology-driven scene interpretation. Research on modeling either spatial or domain-dependent concepts (or both) in an ontology, and exploiting such an ontology for interpreting a graphical scene, is closely connected with our work. Haarslev *et al.* (1994) present one of the first works in this area, introducing “spatio-terminological inferences” to mean a three-level view of inference processes combining quantitative, qualitative, and conceptual representations. They use the TBox and ABox of LOOM (Baader *et al.* 1991) and apply spatio-terminological reasoning to parsing visual programming languages. Other works by the same research group use different ontology languages and address different application domains, but remain consistent with the seminal proposal. As an example, Haarslev *et al.* exploit description logic and apply ontological reasoning to sketch-based queries for Geographical Information Systems (Haarslev 1999; Haarslev *et al.* 2002).

¹¹ The XML Metadata Interchange (XMI), <https://www.omg.org/spec/XMI>, accessed on July 2019.

In his recent book “Description Logics in Multimedia Reasoning”, Sikos (2017) presents an integrated and comprehensive analysis of issues relevant to our work, with chapters on spatial description logics, spatial annotations, and reasoning tools.

Forestier *et al.* (2008) and Bannour and Hudelot (2011) present other ontologies for modeling spatial concepts and reasoning on scenes and images. To make a recent example, Guérin *et al.* (2017) exploit one ontology that formalizes the basic concepts of the image processing domain and provides a way to organize and use input and output data in a formal structure, and provide a formal ontological implementation of the comic books domain. This ontology is meant to handle the content of a comic book, to support the automatic extraction of its visual components, and to formalize the semantics of the domain’s codes.

While taking inspiration from works on spatial ontologies, OntoScene needs to model notions like “Classification” and “Interpretation” that allow us to distinguish between the “syntax” of the image, dealt with by the Detector and Classifier modules, and its semantics, devised thanks to ontological reasoning on the domain, along with logical reasoning. Being a JADE MAS, our framework requires the OntoScene Ontology to be compliant with the JADE requirements for ontology management. For these reasons, we could not reuse existing ontologies as they are; moreover, some ontologies were not available to the research community and others were not modeled in OWL, as needed in our work. Nevertheless, we took them into account when modeling the “GeometricRelations” concept.

3 The OntoScene framework: A gentle introduction

3.1 The initial scenario

Viviana is very curious about the prehistoric rock art of Mount Bego and she would like to know how a domain expert would interpret the image shown in Figure 3 according to the most recent archeological findings.

In that image, Viviana can only see a “matrix” in the top right corner, with a kind of filled trapeze overlapping it, and three symbols very similar to each other, made of lines with filled rectangles in the middle, in the center, and the bottom left corner of the image.

Massimiliano, who is good at detecting and classifying symbols from a purely syntactic point of view, explains her that the “matrix” can be classified as a “Reticulum_Class” with 100% confidence, the trapeze along with the rectangle just below it can be classified as a “Dagger_Class”, and the three symbols made of lines with small filled rectangles in the middle can be classified as “Up_Corn_Class”. These classes are drawn from an ontology modeling information about Mount Bego’s petroglyphs.

Viviana is far from being satisfied, since this syntactic classification says nothing about the meaning of symbols and of the scene as a whole. She sends the information provided by Massimiliano to Daniela, who knows many archeologists, and asks her if she can provide a semantic interpretation of the scene.

Daniela contacts Annie and Henry: Annie is very good in associating domain-dependent meaning to symbol classifications. By exploiting the same ontology used by Massimiliano, she can confirm that a symbol classified as a “Reticulum_Class”, when interpreted inside




Fig. 3. A Complex Scene from the book by de Lumley and Echassoux (2009), page 176, Figure 142. This and all the other images taken from that book are reproduced by kind permission of Professors de Lumley and Echassoux.

a rock art artifact from Mount Bego, actually represents a “Reticulum”; in another domain, the “Reticulum_Class” might have been interpreted as “Prison Bars” or “Chess Board”: decoupling the classification from the interpretation fosters reuse and modularity, and the domain ontology is a good means for achieving this aim. A “Dagger_Class” represents a “Dagger” in the Mount Bego rock art domain, and the “Up_Corn_Class” represents a “Corniform”.

The semantics associated by Annie with the classifications devised by Massimiliano is still not enough to interpret the scene: more knowledge and more reasoning are needed. Taking Annie’s interpretation of symbols belonging to the scene into account, Henry reasons about them and their spatial relationships and finally informs Daniela that the dagger and the reticulate at the top of the image identify the “Storm God” inside a pastoral scene, characterized by a group of corniforms (de Lumley and Echassoux 2009). Another possible interpretation could be that the two corniforms in the center of the image, one inside the other, identify the “Bull God”, and the bottom left corniform is a stand-alone symbol, unrelated with the others. However, Henry thinks that the first interpretation is the most likely one.

Viviana is now happy with this explanation: by moving from symbol classification (symbol syntax) to interpretation (symbol semantics), and then combining interpretations into coherent subscenes via domain-dependent rules, her friends helped her understanding the image.

The people involved in this scenario and the way they interact reflect the OntoScene framework that we developed: each person could be suitably associated with an agent or a component in the OntoScene software framework depicted in Figure 4:

- Viviana is an unnamed, generic agent `AgentX` that wants to understand the meaning of a scene depicted inside an input image: she interacts with a software module (Massimiliano) able to detect coherent sub-images, also named “tokens”, inside

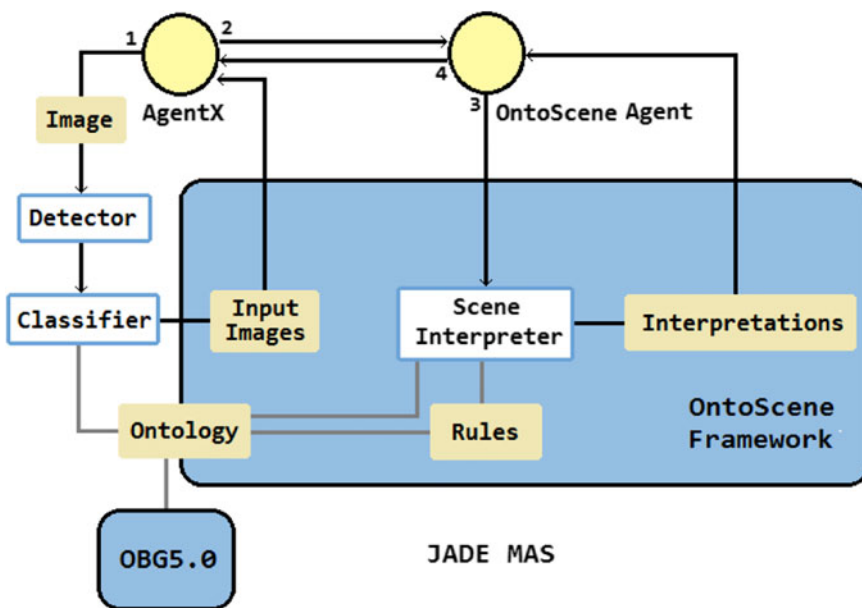


Fig. 4. OntoScene Framework: architecture and data flow.

an image and to classify them, and with another agent (Daniela) that acts as an interface with the domain experts.

- Massimiliano plays the role of token Detector and Classifier, and is able to divide an input image into sub-images. The computed set of sub-images, each one associated with a list of possible classifications, is sent back to AgentX, Viviana in this example.
- Daniela acts as the OntoScene Agent, managing the interactions with Annie and Henry, to provide an interpretation for the image.
- Annie and Henry implement the intelligent engine able to interpret scenes according to the meaning of classified tokens, and to the rules that aggregate such interpretations (also taking spatial relations into account), to provide a semantics of the complex scenes (SceneInterpreter).
- All these agents and components share a common ontology.

To go deeper inside the high-level architecture of OntoScene and the data flow within it, white rectangles in Figure 4 represent system modules, while light yellow (light gray in B/W) rectangles represent either data flowing between them, or data that are used by them. Circles represent agents and the blue (dark gray in B/W) rectangles with rounded corners represent the two platforms involved in the process.

An arrow flowing from A to B tagged with data D, represented as a rectangle on the arrow itself, means that D is generated as an output by A and used as an input by B. An arrow flowing from A to B with no tag means that A generates some output that becomes an input for B (but we do not need to identify it). A gray line between two components means that a “uses/is used” relationship holds between them.

Data managed by OntoScene are:

- Image, the raw input image to be interpreted;




Fig. 5. How the Detector and Classifier modules interact.

- InputImages, the output of the Detector and Classifier representing the input image and the tokens therein, along with their bounding boxes and their classifications, in a symbolic format;
- Prolog Rules, which are set by the domain experts and define how to interpret an image;
- Interpretations, which represent the final output computed by OntoScene;
- Ontology, which represents the application domain, namely the classifications, interpretations, and geometric relationships that are meaningful for the specific image domain and interpretation task; these concepts are used by the Rules (Section 3.5).

3.2 Syntactic pre-processing: Detector and Classifier

The interpretation of the input scene requires that it has been segmented into atomic sub-images (“tokens”) and that one or more classifications have been associated with each of them. To this aim, we assume the availability of a *Detector* and *Classifier*.

We do not enter into the details of how these modules could be designed and implemented, since many libraries and tools for solving the bounding box detection and the classification problems exist and are available to the community. Just to make some examples, the MathWorks Image Processing Toolbox¹² provides algorithms for image processing, analysis, visualization, and segmentation; OpenCV¹³, cross-platform and free for both academic and commercial use, offers 2D segmentation and recognition functionalities suitable for the implementation of both the *Detector* and the *Classifier*, besides many other advanced features; ImageJ¹⁴, written in Java, and Pillow¹⁵, in Python, are other libraries providing edge detection functionalities useful for implementing the *Detector* module.

As far as the classification of images in the rock art domain is concerned, we refer to our previous work within the IndianaMAS project, where *ad hoc* detection and classification algorithms were developed (Briola et al. 2017; Mascardi et al. 2014).

To show how the *Detector* and *Classifier* modules are expected to work, we consider an example. The input image in Figure 5 contains three figures: a rectangle, a triangle, and a circle. The *Detector* identifies the three sub-images and associates them with a bounding box rectangle (BB) representing their position and size within the image. The *Classifier* analyzes the sub-images identified by the *Detector* and assigns the R (rectangle), T (triangle), and C (circle) classifications, consistently with the domain ontology.

¹² <https://www.mathworks.com/products/image.html>, accessed on July 2019.

¹³ <http://opencv.org/>, accessed on July 2019.

¹⁴ <https://imagej.nih.gov/ij/index.html>, accessed on July 2019.

¹⁵ <https://python-pillow.org/>, accessed on July 2019.




Fig. 6. The SceneInterpreter module.

The Classifier is expected to assign multiple classifications to the detected figures, in case of ambiguity. Its output is hence a list of possible classifications for each BB, with an associated confidence in the interval $[0.0, 1.0]$. If there are no doubts about the classification, the list will contain one element only.

3.3 From syntax to semantics: SceneInterpreter

Figure 6 shows the SceneInterpreter, the core module of OntoScene. SceneInterpreter takes an image consisting of a set of tokens in input (we will call this set a “scene”) and returns all its interpretations. It is driven by logical rules that define the possible meanings of each token recognized during the detection and classification stages, and the “well formed” scenes that the framework can recognize and interpret along with their meaning.

A figure classified as a Circle might be interpreted as a Planet in an astronomic domain, as a Face in an emoticon recognition domain, as a Traffic_Light_Element by a self-driving car: the classification as a circle is not enough to correctly interpret a figure in a context made up of other figures. Making the link between the classification and the interpretation levels explicit allows the designer to reuse the classification output and to change the scene interpretation according to the current domain, by only changing the interpretation rules.

As an example, in the rock art domain that provides the case study of this work, a figure classified as an Anthropomorphic_Shape might be interpreted as a Human, a figure classified as a Line_Shape might be either a Sword or a Staff, and a triangle should be interpreted as a mage cap.

The interpretation of an individual token is defined by means of the `interpretation(Cl, ImgInt)` fact that associates the interpretation `ImgInt` with the classification `Cl`. In the rock art example, interpretation facts might look like

```

interpretation('Anthropomorphic_Shape', 'Human').
interpretation('Line_Shape', 'Sword').
interpretation('Line_Shape', 'Staff').
interpretation('Triangle_Shape', 'MageCap').
  
```

Rules that define how to interpret scenes can be presented, in a user-friendly and simplified form, as `rule(SceneInt, ImgList){Cond}`, stating that the scene consisting of sub-images listed in `ImgList` should be interpreted as `SceneInt` based on conditions `Cond`. The conditions involve the interpretations of sub-images in `ImgList` and the spatial relations between/among them.

Table 1. *User-friendly modeling language for scene interpretation rules: boldface symbols are terminals; alphanumeric uppercase strings are defined in the usual way; properties should include at least the geometric binary relations listed in the BNF, but unary properties such as the image color or source, and n-ary properties such as belonging to the same group, could be added*

```

interpretationRule ::= rule(sceneInt , [ imgList ] ){cond}
sceneId ::= uppercase alphanumeric string
imgId ::= uppercase alphanumeric string
interprId ::= uppercase alphanumeric string
sceneInt ::= ` sceneId `
imgList ::= imgId | imgId , imgList
constraint ::= interprId( imgId ) | property( imgList )
disjcond ::= constraint or constraint | constraint or disjcond
cond ::= constraint | ( disjcond ) | constraint ; cond
property ::= horizontal | vertical | diagonal | disjoint | ...

```




Fig. 7. Input scene: first example.

Domain experts may use the user-friendly syntax, whose BNF is presented in Table 1, which can be automatically translated into standard Prolog¹⁶.

As an example, the first rule below can be read as “if token X has been interpreted as a human figure, and if token Y has been interpreted as a sword, and if X and Y are positioned horizontally, then they form a scene representing a Warrior”. The second rule is similar, but states when two tokens represent a Shepherd.

```

rule(`Warrior`, [X,Y]) {
  Human(X);
  Sword(Y);
  horizontal(X,Y);
}
rule(`Shepherd`, [X,Y]) {
  Human(X);
  Staff(Y);
  horizontal(X,Y);
}

```

Let us suppose that the Classifier has classified the leftmost sub-image in Figure 7 as an Anthropomorphic_Shape and the rightmost as a Line_Shape, and the rules above have been loaded into the SceneInterpreter module. Let us also assume that the horizontal geometric relationship holds between the two sub-images. SceneInterpreter generates two interpretations: Warrior(I1) and Shepherd(I2). Interpretation I1 is generated when the rightmost sub-image is interpreted as a Sword (because of the rule for Warrior), while I2 is generated when it is interpreted as a Staff (because of the rule for Shepherd).

¹⁶ The translation has not been implemented so far, but the automatic translation rules are easy to devise, with “;” translated into “,”, **or** translated into “;”, testing of geometric properties translated into `jpl_call` with the property to be tested as argument.




Fig. 8. Input scene: second example.

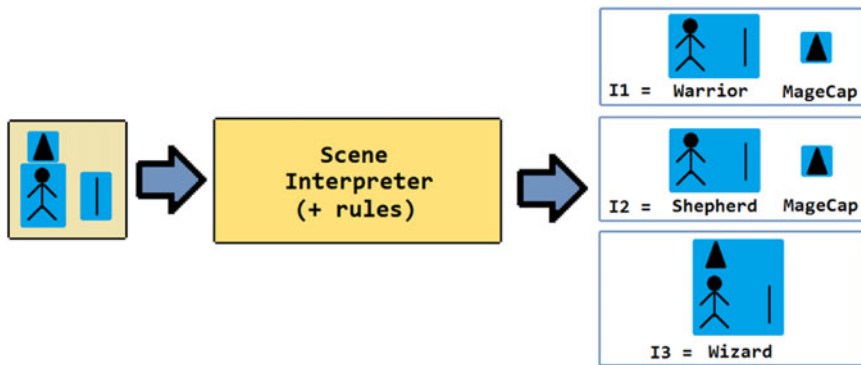


Fig. 9. Input scene: third example.

Figure 8 shows the interpretations of the same scene shown in Figure 7 where a triangular shape has been added on top of the human figure. SceneInterpreter always tries to aggregate as many tokens as possible, but since there are no rules involving the mage cap together with the other elements of the figure, the computed interpretations are those output before, where the triangle is interpreted as a “stand-alone” element.

If another rule were available,

```
rule('Wizard', [X, Y, Z]) {
    Human(X);
    MageCap(Y);
    Staff(Z);
    vertical(X, Y);
    horizontal(X, Z);
}
```

stating that a wizard is a human figure with a magician’s hat on top and a stick placed horizontally, then the SceneInterpreter output would be the one shown in Figure 9.

3.4 Making OntoScene functionalities available to JADE: The OntoSceneAgent

OntoScene has been designed to be a component able to offer the interpretation service, and to be naturally integrated within a JADE MAS. The steps required to perform the integration in a JADE MAS are:

- to integrate the ontology used in the MAS with the OntoScene Ontology in order to allow all agents to be aware of the input and output concepts used within the framework and allow their exchange via JADE messages;

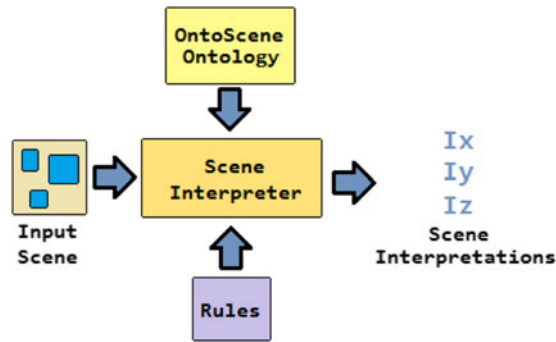


Fig. 10. The OntoScene Ontology in the OntoScene context.

- to add a new JADE action representing the interpretation of a scene (`InterpretAction`): we achieved both these two steps thanks to the OBJ5.0 framework (Briola et al. 2018);
- to implement an agent acting as an interface between the other agents and OntoScene; this agent (the `OntoSceneAgent`) waits for an agent *A* to send a request to perform the action `InterpretAction`, with an input scene, calls the `SceneInterpreter` module on it, and returns the scene interpretations to *A*.

Since this issue is not central to the paper, which focuses on the implementation of the OntoScene framework, we do not expand it further.

3.5 The OntoScene Ontology

To formalize the OntoScene domain and make interoperability among the many modules involved in the framework possible (Figure 10), an ontology called OntoScene Ontology has been designed and implemented.

The OntoScene Ontology is aimed at ensuring modularity and domain independence: the user can extend it by adding more domain concepts from existing or new ontologies. In fact, concepts such as `Classification` and `Interpretation`, which characterize the ontology (see Section 4.1 for more details) are necessarily domain-specific: by changing the domain ontology that extends the OntoScene Ontology, and consistently changing the interpretation rules, the user can modify the application domain while leaving the OntoScene core functionalities unchanged.

3.6 Back to the initial scenario

Thanks to the components mentioned in the previous sections, we can obtain the bounding boxes shown in Figure 11 and the interpretations, represented in a way that should be intuitive enough and that will be explained in details in Section 5, below:

```

I1 = [Storm_God(Reticulum-0, Dagger-1),
Group_Of_Corniforms(Corniform-2, Corniform-3, Corniform-4)].
  
```

```

I2 = [Storm_God(Reticulum-0, Dagger-1),
Bull_God(Corniform-2, Corniform-3), Corniform-4].
  
```




Fig. 11. The Scene from the book by [de Lumley and Echassoux \(2009\)](#) with detected bounding boxes.

Given that `bb(X, Y, W, H)` states the X and Y coordinates of the top- and leftmost corner of the bounding box plus its Width and Height, this is the actual result we get by running OntoScene on the input

```
image(0, bb(161, 12, 165, 167), [class('Reticulum_Class', 1.0)]).
image(1, bb(257, 68, 109, 281), [class('Dagger_Class', 1.0)] ).
image(2, bb(86, 323, 162, 129), [class('Up_Corn_Class', 1.0)] ).
image(3, bb(107, 337, 181, 162), [class('Up_Corn_Class', 1.0)] ).
image(4, bb(3, 506, 144, 23), [class('Up_Corn_Class', 1.0)] ).
```

and includes the correct interpretation I1 provided by Henry de Lumley and Annie Echassoux, two archeologists who spent their life on rock art interpretation, in the book from which the image is taken.

4 Modeling and implementing domain and spatial knowledge

4.1 Domain knowledge

The OntoScene Ontology imports the JADE template ontology, needed to let the ontology be directly usable by JADE, as described in Section 2.1. It contains all the concepts that SceneInterpreter uses during a scene interpretation and is designed to be extended with an existing domain ontology to integrate SceneInterpreter within a MAS in a transparent way. The classes provided by the OntoScene Ontology are shown in Figure 12.

Point. The Point class contains two single *float* properties X and Y.

BoundingBox. The BoundingBox class, abbreviated as BB, represents the rectangle that bounds a single image.

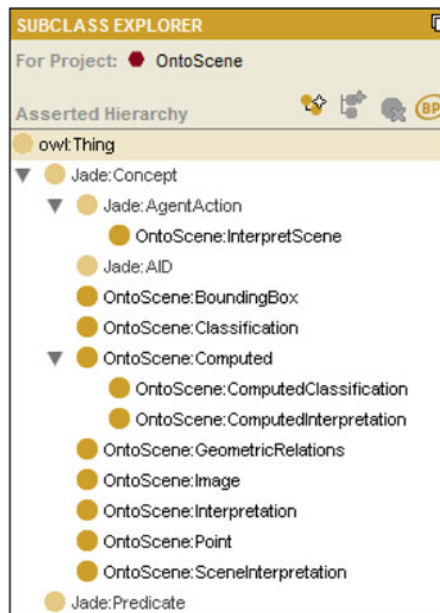


Fig. 12. The OntoScene Ontology.

ComputedClassification. The `ComputedClassification` class represents a classification computed by the `Classifier` along with its confidence. It contains the single properties `identifiedClassification` with range `Classification` and confidence with range *float*.

ComputedInterpretation. The `ComputedInterpretation` class represents an interpretation computed by `SceneInterpreter` with the associated confidence and its size, namely how many input images have been aggregated. It contains the single properties `identifiedInterpretation` with range `Interpretation`, confidence with range *float* and size, with range *int*.

Classification and Interpretation. `Classification` and `Interpretation` are two classes without any property and their meaning is the intuitive one. To allow `SceneInterpreter` to interpret an input scene, some classes from the domain ontology must necessarily extend these two classes with domain-specific classifications and interpretations.

GR. The `GR` class is used as a container for methods representing geometric relationships, to be called within the body of rules through predicates offered by the JPL Library. `SceneInterpreter` uses an internal class called `GeometricRelationsImpl` with the implementation of those methods that we used to test the program. More sophisticated implementations can be used instead of the ones we provide: the Java Method Mapper panel of OBG5.0 allows the developer to create methods under the `GR` class and export their interface, in order to be implemented.



Fig. 13. The Image class. The “multiple” attribute associated with classifications, interpretations, and subParts, means “list of”.

Image. The Image class represents a basic or composite scene. It contains a single id property of type *int* that acts as an identifier, a single boundingBox property of type *BoundingBox* for the BB, a multiple classification property of type *ComputedClassification* listing all the classifications assigned by the Classifier to the image in the scene, a multiple interpretation property of type *ComputedInterpretation* including the interpretations computed by SceneInterpreter and a multiple subParts property of type *Image* that contains all the sub-images that form the image, as shown in Figure 13.

The Image class is the main data structure used by SceneInterpreter to keep track of the relationship between Prolog scenes represented as Prolog facts, and Java scenes represented as instances of the Java Image class. Each time a new node (namely, a new scene) is added to the scene graph, the corresponding Image instance is also created inside it: there is a one-to-one association between each node in the scene graph and an Image instance. In the sequel, we will usually use *image* and *sub-image* when we refer to data representations on the Java side, and *scene* and *sub-scene* when we refer to the Prolog side.

In order to work properly, SceneInterpreter expects input images with these features:

- id and boundingBox fields instantiated;
- classifications instantiated with a list of one or more classifications;
- empty interpretations list;
- empty subParts list.

The association between classifications and interpretations is computed by the Prolog engine via the *interpretation/2* predicate introduced in Section 3.3.

After the creation, via the aggregation rules, of a composite scene in Prolog, SceneInterpreter creates a new Image object that corresponds to the new scene and has these features:

- id field instantiated with a new unique identifier;
- boundingBox obtained by merging the BBs of the subscenes;
- empty classifications list, as only basic scenes have a classification;
- interpretations list containing the computed interpretations;
- subParts instantiated with the list of the sub-images.

SceneInterpretation. The SceneInterpretation class represents an interpretation of the input scene. It contains a *composedBy* property of type *Image* that contains all the

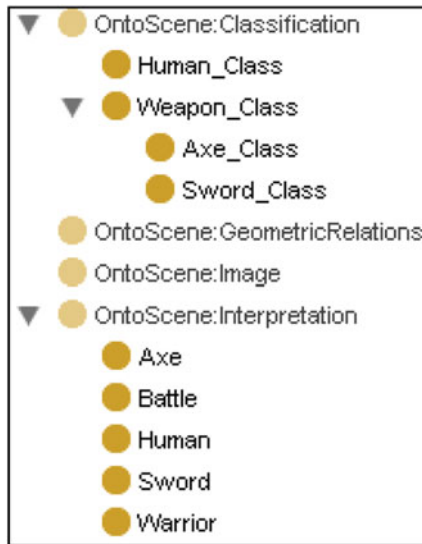


Fig. 14. Domain-dependent concepts that extend OntoScene classes.




Fig. 15. The files generated by OBG5.0.

images of the interpretation in the format presented above, corresponding to the scenes that can coexist.

The agent that, upon reception of an InterpretScene action presented below, is required to provide a scene interpretation, returns a SceneInterpretation list.

InterpretScene action. The InterpretScene class extends the JADE AgentAction class and represents the action of requesting the interpretation of an input scene. It contains a multiple property inputImages of type Image representing the images in the input scene and two *boolean* properties, *distinct* and *filtered*, which refer to the interpretation mode. When *distinct* mode is selected, all the scenes in the final list of SceneInterpretation must be distinct Java objects, in order to obtain a readable and writable data structure. When *filtered* mode is on, only filtered interpretations are returned.

An example ontology: Battle. The Battle ontology models a simplified domain that will be used in the next section. Figure 14 shows how the Classification and Interpretation classes of OntoScene can be sub-classed by classes characterizing the Battle domain, where armed warriors fight using swords or axes. The Java files generated by OBG5.0 are shown in Figure 15.




Fig. 16. Graphical representation of the formula used by `relNear`.

4.2 Spatial knowledge

To interpret scenes with `SceneInterpreter`, the user must identify the required geometric relationships and must create methods in the `GR` class of the `OntoScene` Ontology to represent them. If the user has no special requirements, (s)he can use the `GRImpl` we provide with the framework. Implementing geometric relationships is not easy, because different domains may need different relationships. An exception are topological relationships (disjoint, overlap, etc.) for which known mathematical formalisms exist. We used the `JTS` library to implement the following ones:

```

Boolean horizontal (BB bb1, BB bb2, String pos)
Boolean horizontal (BB bb1, BB bb2)
Boolean vertical (BB bb1, BB bb2, String pos)
Boolean vertical (BB bb1, BB bb2)
Boolean diagonal (BB bb1, BB bb2, String pos)
Boolean diagonal (BB bb1, BB bb2)
Boolean disjoint (BB bb1, BB bb2)
Boolean overlap (BB bb1, BB bb2)
Boolean contains (BB bb1, BB bb2)
Boolean absNear (BB bb1, BB bb2, float th) /* absolute proximity */
Boolean relNear (BB bb1, BB bb2, float th) /* relative proximity */
Boolean absGroup (List <BB> bbs, float th) /* group, using absNear */
Boolean relGroup (List <BB> bbs, float th) /* group, using relNear */

```

Horizontal, vertical, and diagonal relationships. The parameters of these methods are two BBs and – optionally – a string indicating the position that `bb1` must have w.r.t. `bb2`. The position may be *right* or *left* for horizontal, *up* or *down* for vertical, and *se*, *sw*, *ne*, *nw* for diagonal. For example, `diagonal(bbx, bby, ne)` is true if `bbx` is positioned northeast w.r.t. `bby`.

Topological relationships disjoint, overlap, and contain. These methods take two BBs `bbx` and `bby` in input and answer whether `bbx rel bby` holds. For example, `contains(bbx, bby)` is true if `bbx` contains `bby`.

Absolute proximity AbsNear and relative proximity RelNear. Besides the two BBs, these methods also have a third parameter to state the threshold under which the two BBs are considered “close”. This threshold therefore defines the proximity semantics.

In `absNear`, the threshold indicates an absolute value expressed in an arbitrary measure unit determined by the domain expert such as pixels, centimeters. For example, assuming pixels as the measure unit, `absNear(bbx, bby, 10.0)` is true if the absolute distance between the edges of `bbx` and `bby` is less than 10px.

In `relNear`, the threshold indicates a relative value between 0 and 1.0. This allows us to define “proximity” in a way robust to the image scaling. For example, `relNear (bbx, bby, 0.2)` is true if $X \leq 0.2$, where X is the value of some expression that the user can define. The one we implemented is explained in Figure 16: we compute `JTSDist`, namely the distance between `bbx` and `bby` computed by `JTS`, we merge `bbx` and `bby` into `mbb`, we compute `Diagonal`, namely the length of `mbb` diagonal. X is `JTSDist/Diagonal`. If both bounding boxes are scaled by a factor F , `relNear (bbx*F, bby*F, 0.2)` is the same as `relNear (bbx, bby, 0.2)`, making the definition invariant w.r.t. scaling.

Finally, the `absGroup` and `relGroup` methods compute the “neighborhood” relationship on a list of BBs using `absNear` and `relNear`, respectively.

5 SceneInterpreter

The `Detector` and `Classifier` modules work on raw images and produce an “input image” consisting of bounding boxes associated with possibly many classifications of their content, drawn from an ontology, along with a confidence on that classification. `SceneInterpreter` takes this classified “input image” as input and transforms it into a set of “basic scenes”, namely triples consisting of (image, classification, and interpretation).

For each input image, `SceneInterpreter` creates as many basic scenes from the (classification, interpretation) pairs as it can. For example, if a sub-image `Img1` has been classified by the `Classifier` module as `C1` or `C2`, and `C1` has `I11` and `I12` as possible interpretations, while `C2` can only be interpreted as the `I21`, three basic scenes are generated:

```
basic_scene (Img1, C1, I11) .
basic_scene (Img1, C1, I12) .
basic_scene (Img1, C2, I21) .
```

The scene interpretation rules that drive `SceneInterpreter` define how to aggregate the elements in a scene, be they atomic sub-images or scenes, depending on the geometric relationships holding among them. We name them *aggregation rules* in the remainder. Aggregation rules have been also called “scene interpretation rules” in the paper; in this section, we prefer to use “aggregation” to clearly differentiate them from the `interpretation` predicate that will be presented in Section 5.1, which associates an interpretation to a basic image, based on its classification. A composite scene is a scene created by the aggregation of other scenes, which may be in turn basic or composite ones. We talk about scene, without further distinction, when it is not necessary to distinguish whether the scene is a basic or a composite one. `SceneInterpreter` generates a scene graph representing all the scenes that can be derived by applying the aggregation rules to the basic scenes generated from an input image.

As an example, the figure in Table 2, left, shows a scene graph resulting from an input scene containing five different sub-images: they have been transformed into five basic scenes (`BS1`, `BS2`, `BS3`, `BS4`, and `BS5`), and then into [composite] scenes thanks to the available aggregation rules. For example in this case, by applying some aggregation rule, `BS1` and `BS2` can be aggregated into `CS1`. `BS2`, `BS3`, and `BS4` can be aggregated into `CS2`, and so on. We point out that `BS2` was used by an aggregation rule to form `CS1`,

Table 2. A scene graph created from an input image with five sub-images, plus the generated interpretations.

Input image	Interpretations
	<p>I1 = [BS1, BS2, BS3, BS4, BS5] I2 = [CS1, BS3, BS4, BS5] I3 = [CS1, BS3, CS3] I4 = [BS1, CS2, BS5] I5 = [BS1, CS4]</p>

and by another to form CS2. In the same way, BS4 can be used to form both CS2 and CS3. BS2 and BS4 are called shared scenes. The scene graph is oriented (from top to bottom) and acyclic. A top node, or top scene, is a node with no incoming edges. In the figure in Table 2, left, CS1 and CS4 are top nodes.

SceneInterpreter core functionalities have been implemented in Prolog. For efficiency issues, however, geometric relationships have been implemented in Java and are called by Prolog through the JPL Library introduced in Section 2.1.

The steps to be performed to set up SceneInterpreter and to interpret an input image are the following:

1. define the aggregation rules in Prolog (done only once);
2. initialize the Java SceneInterpreter module;
3. select the aggregation rules;
4. load a scene composed of a list of images plus their classification (the output of the Detector and Classifier modules), serializing them into basic scenes;
5. apply aggregation rules to create composite scenes and generate the scene graph;
6. generate all the interpretations by calling the `knuth_algo_x` predicate on the scene graph;
7. filter out interpretations that can be derived from others (optional) and provide the final sorted result.

The steps from 4 to 7 are discussed in Sections from 5.1, 5.2, 5.3, and 5.4, respectively.

5.1 Serializing images in basic scenes

To allow SceneInterpreter to serialize input images into Prolog scenes, associations between classifications and domain interpretations created under the Classification and Interpretation ontology classes must be provided. The predicate that OntoScene offers to this aim is

```
interpretation/2
interpretation(Class, Inter) .
```

whose meaning is that a picture classified as `Class` can be interpreted as `Inter`. For the classification and interpretation within the Battle domain, we defined the following facts:

```
interpretation('Human_Class', 'Human').
interpretation('Sword_Class', 'Sword').
interpretation('Axe_Class', 'Axe').
```

The `Human_Class` classification can be directly interpreted as `Human`, the `Sword_Class` as `Sword`, and the `Axe_Class` as `Axe`. During the image serialization, these facts are used by `SceneInterpreter` to create the basic scenes.

A predicate called `scenes/6` is used to represent basic and composite scenes in Prolog. The signature of the predicate is the following:

```
scenes/6.
scenes(ID, BB, Class, Inter, Conf, SS).
```

- `ID` is the identifier that Prolog uses to identify scenes¹⁷;
- `BB` is the reference to the Java object representing the `BoundingBox` of the image in the input scene;
- `Class` is the classification of the image from which this scene comes from. The field is instantiated in basic scenes and is empty in composite scenes;
- `Inter` is the interpretation of the scene. For basic scenes the variable is instantiated by calling the `interpretation/2` predicate, while for composite scenes the value to associate with the variable is computed by applying the aggregation rules;
- `Conf` is the confidence of the interpretation associated with the scene. For basic scenes whose confidence in the classification is `C`, `Conf` is computed as $C * (1.0 / \text{Count})$, where `Count` is the number of interpretations associated with the scene. For composite scenes, $\text{Conf} = (\text{Conf1} + \text{Conf2} + \dots + \text{ConfN}) / N$ where `N` is the number of aggregated scenes, and `ConfX` is the confidence of `X` scene;
- `SS` stands for `SubScenes` and is the list of the IDs of the basic scenes belonging to the scene.

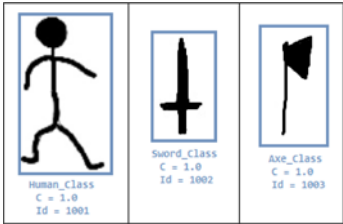
The serialization algorithm is, in pseudocode, the following:

```
InputScene S;
For (Image img: S.getImages ())
  For (Classification class: img.getClassifications ())
    For (Interpretation inter: interpretation (class, inter))
      Assert (scenes (ID, BB, class, inter, Conf, SS))
```

That is, given an input scene `S`, for each sub-image `img` belonging to `S`, for each classification `class` of `img`, for each interpretation `inter` found by calling the Prolog `interpretation/2` predicate, the fact `scene` with suitable arguments is asserted in the Prolog knowledge base, for efficient retrieval. Each individual input image is subdivided into as many basic scenes as the found `(class, inter)` pairs.

¹⁷ The `Id` property of the `Image` class is used by Java and may be different from `ID`.

Table 3. *Input Scene, example 1*

Input Image	Interpretations
	<pre>scene(0, BB1, 'Human_Class', 'Human', 1.0, [0]). scene(1, BB2, 'Sword_Class', 'Sword', 1.0, [1]). scene(2, BB3, 'Axe_Class', 'Axe', 1.0, [2]).</pre>

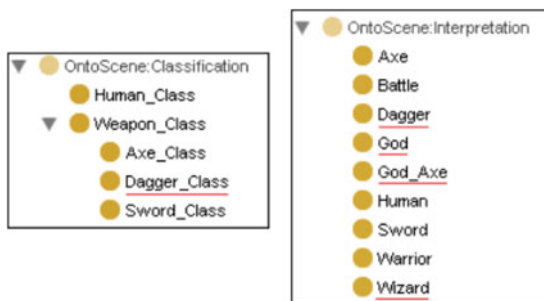


Fig. 17. Extending the Battle ontology with new classifications and interpretations.

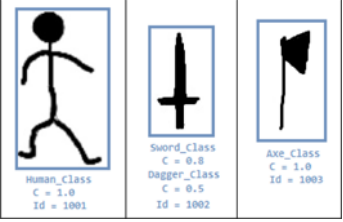



For example, let us suppose that the input scene consists of three sub-images shown in Table 3, classified as Human_Class, Sword_Class, and Axe_Class with maximum confidence. Images are serialized in three scene Prolog facts as shown in the right part of the table.

In the first example, each classification is associated with only one interpretation defined by the domain ontology, but in general there could be a one-to-many relationship. Let us now make the example more complex by adding the Dagger_Class classification and the Dagger, God, God_Axe, and Wizard interpretations (Figure 17). New interpretation facts could be defined as:

```
interpretation('Human_Class', 'God').
interpretation('Human_Class', 'Wizard').
interpretation('Axe_Class', 'God_Axe').
interpretation('Dagger_Class', 'Dagger').
```

In a second example shown in Table 4, the image in the center can be classified into two ways: Sword_Class and Dagger_Class (each having only one interpretation), while the image on the left has one classification Human_Class with three interpretations (Human, God, and Wizard). The image on the right has one classification (Axe_Class) and two interpretations (Axe and God_Axe). The confidence is $1.0 * (1.0/3) = 0.33$ for each interpretation of the left image, is $0.8 * (1.0/1)$ and $0.5 * (1.0/1)$ for the two interpretations of the image in the center, and is $1.0 * (1.0/2)$ for the image on the right.

Table 4. *Input Scene, example 2*

Input Image	Interpretations
 <div style="display: flex; justify-content: space-around;"> <div style="border: 1px solid blue; padding: 2px; width: 30%;">  Human_Class C = 1.0 Id = 1001 </div> <div style="border: 1px solid blue; padding: 2px; width: 30%;">  Sword_Class C = 0.8 Dagger_Class C = 0.5 Id = 1002 </div> <div style="border: 1px solid blue; padding: 2px; width: 30%;">  Axe_Class C = 1.0 Id = 1003 </div> </div>	<pre> scene(0, BB1, 'Human_Class', 'Human', 0.33, [0]). scene(0, BB1, 'Human_Class', 'God', 0.33, [0]). scene(0, BB1, 'Human_Class', 'Wizard', 0.33, [0]). scene(1, BB2, 'Sword_Class', 'Sword', 0.8, [1]). scene(1, BB2, 'Dagger_Class', 'Dagger', 0.5, [1]). scene(2, BB3, 'Axe_Class', 'Axe', 0.5, [2]). scene(2, BB3, 'Axe_Class', 'God_Axe', 0.5, [2]). </pre>

5.2 Applying aggregation rules for composite scenes and updating the scene graph

After defining the `interpretation/2` predicate for the basic scenes, it is necessary to create aggregation rules for composite scenes. We use the `predicate rules/2`, stating which scenes should be aggregated, which geometric relationships between their BBs should hold, and computing a list of scene facts that `SceneInterpreter` uses to generate (possibly) a new composite scene, with interpretation `Inter`.

The clauses for the rule predicate, which are semi-automatically compiled into Prolog from the user-friendly modeling language presented in Table 1, follow this pattern:

```

rule(Inter, Scenes) : -
% Part 1: Selects the scenes to be aggregated in the Scenes list
% Part 2: Computes geometric relationships

```

These rules convey the very same meaning and structure as those presented in Section 3.3; they are less readable since they use the concrete Prolog syntax and JPL calls to spatially related methods based on JTS. For the sake of clarity, we will abuse Prolog notation using `ImgInt(X)` to mean that token `X` has been interpreted as `ImgInt`. The (manual) process for compiling the user-friendly modeling language into Prolog is not optimized: this can be noticed for example in the usage of `append` in Table 5, which could be avoided using unification instead. While losing in elegance of the resulting code, the naive manual compilation produced rules which follow the same pattern and gave useful hints on how they implement the automatic compilation, which will be addressed as a close future work.

Two utility predicates used inside rule clauses are

```

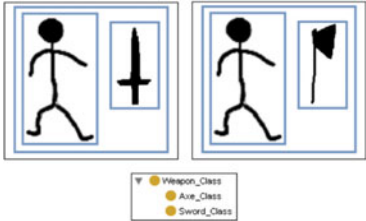
relations/1
relations(GR).

and

subclass_of/2
subclass_of(Class, SubClass).

```

Table 5. *Warrior Scene (Human + Weapon). Image and rule*

Example image	Scene interpretation rule
	<pre>rule('Warrior', Scenes) :- scene(ID1, BB1, Class1, 'Human', Conf1, SS1), subclass_of('Weapon_Class', Class), interpretation(Class, Weapon), scene(ID2, BB2, Class2, Weapon, Conf2, SS2), append([scene(ID1, BB1, Class1, 'Human', Conf1, SS1)], [scene(ID2, BB2, Class2, Weapon, Conf2, SS2)]), Scenes), relations(GR), jpl_call(GR, horizontal, [BB1, BB2], @(true)), jpl_call(GR, near, [BB1, BB2, 2.0], @(true)).</pre>

relations(GR) unifies GR with a reference to the implementation of the interface for the geometric relations, instantiated during the OntoScene configuration stage via a call to `jpl_new/3`. In our code, the assertion of the relations(GR) predicate is achieved via

```
assert_relations :-
  jpl_new('onto_impl.GeometricRelationsImpl', [], GR),
  assert(relations(GR)).
```

Other OntoScene users might use our implementation of geometric relations, provided via the `'onto_impl.GeometricRelationsImpl'` interface, or develop a new one.

The `subclass_of(Class, SubClass)` is a predicate exported with OBG5.0: it allows scenes to be analyzed by exploring hierarchies of classes in the ontology, in particular those below the `Classification` and `Interpretation` classes.

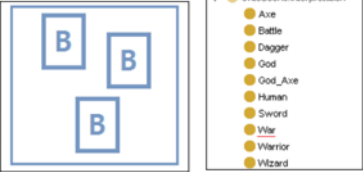
Each scene generated by applying one aggregation rule is asserted as a node of the scene graph which is modeled via the `image_graph(G)` fact, and which is updated any time a new scene interpretation is computed for a given image, reaching at the end the structure exemplified in Table 2.

In the sequel, we provide some examples of aggregation (scene interpretation): `near` is used as an abbreviation for `absNear` and lengths are expressed in pixels.

Example 1: Warrior Scene (Human + Weapon). A generic `Warrior` scene can be defined as a combination of a `Human` scene and a basic scene classified as `X`, where `X` is a subclass of `Weapon_Class` in the ontology (Table 5).

A composite scene can be defined by other composite scenes. For example, if we want to define a `Battle` scene as a combination of two composite `Warrior` scenes, a rule could be defined to check that two `Warrior` scenes have been detected in the image, and that they are close enough. In general, the user can implement rule in any way, using all the expressive power of Prolog and creating auxiliary predicates for designing and implementing more complex rules. The rules presented so far only aggregate two scenes at a time, but of course it is possible to select a larger number. For example, a scene of `War` could be formed by an arbitrary number of `Battle` scenes close to each other, as shown in the next paragraph.

Table 6. War scene (group of Battles)

Example image	Scene interpretation rule
	<pre>rule('War', Scenes) :- findall(scene(ID, BB, Class, 'Battle', Conf, SS), scene(ID, BB, Class, 'Battle', Conf, SS), Battles), sublist(Battles, Scenes), relations(GR), jpl_call(GR, group, [BBs, 10.0], @(true)).</pre>

Example 2: War scene (group of Battle). Figure in Table 6 shows a War scene consisting of three Battle scenes, close to each other. The rule implementation could be the one on the right of the table, which looks for all the asserted Battle scenes and nondeterministically selects some of them using the `sublist/2` predicate. Finally, it checks that those scenes are close enough to form a group (`jpl_call(GR, group, [BBs, 10.0], @(true))`).

5.3 Computing all the possible interpretations

The main functionality of `SceneInterpreter` consists of analyzing all the nodes in the scene graph to determine which of them can coexist in an interpretation (which has to contain all the basic scenes). Two nodes can coexist in the same interpretation if and only if they do not share any basic scene. For example, in the figure in Table 2, node CS1 and node CS2 cannot coexist in an interpretation because they share BS2.

This “coexistence check” resorts to the NP-complete exact cover problem (Karp 1972). Let X be the set of the basic scenes computed, and asserted, in the way discussed in Section 5.1. Each node in the scene graph identifies a subset of X : the scene graph is a collection S of subsets of a set X . By definition, an exact cover of X is a subcollection S^* of S that satisfies two conditions:

1. The intersection of any two distinct subsets in S^* is empty, that is, the subsets in S^* are pairwise disjoint. In other words, each element in X is contained in at most one subset in S^* .
2. The union of the subsets in S^* is X , that is, the subsets in S^* cover X . In other words, each element in X is contained in at least one subset in S^* .

A subcollection S^* satisfying the two properties above is indeed what we name a *scene interpretation*. `SceneInterpreter` implements Donald Knuth’s *Algorithm X* for the exact cover problem (Knuth 2000). Algorithm X is a recursive, nondeterministic, depth-first, backtracking algorithm: the ideal algorithm for Prolog!

If we disregard the code for managing matrices (an `update_matrix` predicate is needed, whose code is not shown), the Algorithm X’ Prolog implementation is 14 lines long, excluding comments.

The exact cover problem is represented in Algorithm X using a matrix A consisting of 0s and 1s. The goal is to select a subset of the rows so that the digit 1 appears in each

Table 7. *Donald Knuth's Algorithm X implementation in Prolog*

```

% if the matrix is empty, terminate by unifying the last argument
% with the accumulator
knuth_algo_x([], _, _, Solution, Solution) :- !.

% otherwise
knuth_algo_x(M, Nodes, NumNodes, AccSolution, Solution) :-

    % deterministically select one column containing as few 1s
    % as possible
    transpose(M, TM),
    get_col_with_less_ones(TM, NumNodes, 0, 0, ColIndex, ColCount),
    ColCount > 0,

    % non deterministically select one row with the selected
    % column equal to 1
    member(Row, M),
    nth0(ColIndex, Row, 1, _),

    % update the partial solution
    nth0(Index, M, Row, _),
    nth0(Index, Nodes, SolutionNode, _),
    append(AccSolution, [SolutionNode], NewAccSolution),

    % update the matrix
    findall(I, nth0(I, Row, 1, _), ColumnsToRemove),
    update_matrix(M, ColumnsToRemove, NewM, RemovedRowsIndexes),

    % remove the nodes that were associated with the removed rows
    findall(N, (member(N, Nodes), nth0(I1, Nodes, N, _),
    \+ member(I1, RemovedRowsIndexes)), NewNodes),

    % recursively call the algorithm on the reduced matrix and nodes
    knuth_algo_x(NewM, NewNodes, NumNodes, NewAccSolution, Solution).

```

column exactly once. Table 7 shows the Prolog code for the algorithm, implemented by the `knuth_algo_x` predicate:

```

knuth_algo_x/5.
knuth_algo_x(M, Nodes, NumNodes, AccSolution, Solution).

```

- `M` represents the matrix associated with the collection `S` of subsets of `X` (which, in turn, is associated with the scene graph stored via the `image_graph(G)` fact); it is represented in a standard way as a list of lists, making it possible to exploit the `transpose/2` predicate offered by the SWI-Prolog CLP(FD) library for Constraint Logic Programming over Finite Domains¹⁸.
- `Nodes` is the list of nodes in the scene graph.
- `NumNodes` is the number of nodes in the scene graph.

¹⁸ https://www.swi-prolog.org/pldoc/doc/_SWI_/library/clp/clpfd.pl, accessed on July 2019.

- `AccSolution` is the accumulator argument.
- `Solution` is unified with the solution, when the algorithm terminates.

The nondeterministic choice of the row via the `member(Row, M)` goal allows the algorithm to “clone” itself into independent subalgorithms which work on a reduced version of the matrix `M`. Searching the state space is of course left to the Prolog interpreter.

Each set of nodes in the graph scene which is an exact cover of the basic scenes is an interpretation of the input scene. The possible interpretations of the figure in Table 2, left, are reported in the table right side.

The set `[BS1, BS2, BS3]` is not an interpretation because it does not contain all the input images (`BS4` and `BS5` are missing) and `[CS1, CS4]` is not correct as well because both `CS1` and `CS4` share the same scene `BS2` (and hence cannot coexist).

5.4 Filtering, sorting, and returning interpretations

Usually, one input image generates many scene interpretations, some of which can be derived from others by substituting one aggregated scene with the subscenes which form it. `SceneInterpreter` can filter out interpretations that can be derived by others in this way. In the example above, `I1` and `I2` can be derived from `I3` and can be filtered out: if we substitute `CS1` with its children `BS1` and `BS2`, and (resp. or) `CS3` with `BS4` and `BS5`, we obtain `I1` (resp. `I2`).

Each interpretation is checked against the others computed so far, to avoid duplicates due to order of nodes in the interpretation, and is associated with a weight computed as the sum of the squares of the aggregated scenes lengths. As an example, the weight of the following interpretations is 8 and 4, respectively.

```
I1 = [W1(Human-0, Sword-1), W2(Human-3, Sword-2)]. % Weight = 8
I2 = [Human-0, Sword-1, Human-3, Sword-2].       % Weight = 4
```

Interpretations are sorted in decreasing weight order, from the one which aggregates more scenes together to the one where less aggregation rules have been exploited. In the example above, `I1` “aggregates more” than `I2` and comes before `I2` in the list of computed interpretations, but both are returned.

5.5 SceneInterpreter at work

In this section, we show further examples in the Battle domain, each coming with an informal description and the interpretations that the Prolog interpreter generates when the `generateAllInterpretations` and `generateFilteredInterpretations` predicates are called. We consider images whose possible classifications are `Human_Class`, `Sword_Class`, and `Axe_Class`, with maximum confidence. For each classification, we assume that only one interpretation exists:

```
interpretation('Human_Class', 'Human').
interpretation('Sword_Class', 'Sword').
interpretation('Axe_Class', 'Axe').
```

In Table 8, the basic scenes generated for the images that will be used in the examples are reported.

Table 8. Basic scenes used in the next examples

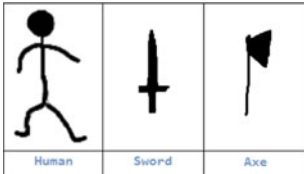
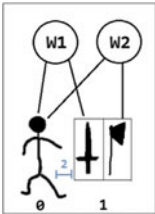
Input Image	Basic Scenes
	<pre>scene(ID1, BB1, 'Human_Class', 'Human', 1.0, [0]). scene(ID2, BB2, 'Sword_Class', 'Sword', 1.0, [1]). scene(ID3, BB3, 'Axe_Class', 'Axe', 1.0, [2]).</pre>

Table 9. Example 1 of a complex scene with its interpretations

Input image	Interpretations
	<pre>generateAllInterpretations (4): I1 = [W1 (Human-0, Sword-1)]. I2 = [W2 (Human-0, Axe-1)]. I3 = [Human-0, Sword-1]. I4 = [Human-0, Axe-1]. GenerateFilteredInterpretations (2): I1 = [W1 (Human-0, Sword-1)]. I2 = [W2 (Human-0, Axe-1)].</pre>

In the next examples, for each scene, we show the scene graph (generated by calling the `applyRules` method) and the generated interpretations. We consider the following composite scenes:

Warrior = Human + Weapon (Sword or Axe)

with distance between the BB of Human and the BB of Weapon ≤ 2 px.

Battle = Warrior + Warrior

with distance between the BBs ≤ 5 px.

Example scene 1. Table 9 shows a Human (ID = 0) close to another figure (ID = 1) that can be classified as `Sword_Class` and `Axe_Class`, and hence interpreted as `Sword` and `Axe`. The scene graph generated by `applyRules` contains two Warriors, W1 and W2. The generated interpretations are reported on the right of the table.

Example scene 2. Table 10 shows four Humans and four Weapons. Each Human is close enough to the Weapon at its right to be interpreted as a Warrior (W1, W2, W3, W4), and each Warrior is close enough to the adjacent Warrior to be considered as a Battle (B1, B2, B3). The first five generated interpretations, on a total of 29 ones, are reported on the right of the table.

Table 10. Example 2 of a complex scene with its interpretations

Input Image	Interpretations
	<pre> generateAllInterpretations (29): I1 = [B1 (W1, W2), B3 (W3, W4)]. I2 = [W1 (Human-0, Sword-1), B2 (W2, W3), W4 (Human-6, Axe-7)]. I3 = [W1 (Human-0, Sword-1), W2 (Human-2, Sword-3), B3 (W3, W4)]. I4 = [B1 (W1, W2), W3 (Human-4, Sword-5), W4 (Human-6, Axe-7)]. I5 = [Human-0, Sword-1, B2 (W2, W3), W4 (Human-6, Axe-7)]. (24 more interpretations are generated). GenerateFilteredInterpretations (2): I1 = [B1 (W1, W2), B3 (W3, W4)]. I2 = [W1 (Human-0, Sword-1), B2 (W2, W3), W4 (Human-6, Axe-7)]. </pre>

Table 11. Example 3 of a complex scene with its interpretations

Input image	Interpretations
	<pre> generateAllInterpretations (3): I1 = [B1 (W1, W2)]. I2 = [W1 (Human-0, Sword-1), W2 (Human-3, Sword-2)]. I3 = [Human-0, Sword-1, Human-3, Sword-2]. GenerateFilteredInterpretations (1): I1 = [B1 (W1, W2)]. </pre>

Example scene 3. Table 11 shows two Humans on the right and on the left of the picture, both close to the two Swords in the center. Each Human can only be associated with the Sword that is closest to him (Human 0 cannot be associated with Sword 2 and the same for 3 and 1). Hence, the only possible interpretations are two Warriors W1 and W2 and one Battle B1. The generated interpretations are reported on the right of the table.

6 Case study: Interpreting scenes from the rock art domain

In this section, we present OntoScene at work. The domain where we experimented it is the one introduced in Section 3.1: Mount Bego’s prehistoric rock art.

6.1 Studies by Clarence Bicknell and Henry de Lumley

Archeologists and historians look at the area around Mount Bego as an incredibly valuable source of knowledge, due to the up to 40,000 figurative petroglyphs and 60,000 nonfigurative petroglyphs scattered over a large area at an altitude of 2000 m to 2700 m.

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.

The explorer who first realized the importance of Mount Bego carvings was Clarence Bicknell, who, at the turn of the 20th century, created an important catalog of most of the petroglyphs in Mount Bego (Bicknell 1913).

Many years after Bicknell's campaigns, several teams led by Henry de Lumley have been surveying and mapping this archeological area starting from 1967 (Bianchi 2011; de Lumley and Echassoux 2009).

The University of Genova owns a collection of 16,000 drawings and reliefs made by Clarence Bicknell between 1898 and 1910, in his campaigns on Mount Bego. Bicknell's Legacy also includes nine notebooks, filled with notes in Victorian English, mostly unpublished. The publication on the web of about 350 images from the Bicknell's drawings and reliefs (Rolls 8, 20, 23, available on the Bicknell Legacy website) along with their classification was one of the results of the IndianaMAS research project.

The images used for the experiments presented in this section and in the Appendix come from the Bicknell's Legacy and from the book by de Lumley and Echassoux (2011): we report an identifier under each image to refer to the first (abbreviated into BL, R. for Roll and P. for page) or to the second (abbreviated into DE, P. for page and F. for figure number).

For each type of scene in the dataset, three or four images were manually selected to represent the most frequent recognized patterns. The Detector and Classifier modules were simulated by manually drawing BBs around the sub-images of the scene and assigning them the classifications provided by Dr. Nicoletta Bianchi, who collaborated with us in the IndianaMAS project and in the construction of the Bicknell Legacy website. With her help, we also produced a natural language interpretation rule for Bicknell's images and we translated them in Prolog for each scene type. As far as de Lumley and Echassoux' images are concerned, the natural language interpretation rules are those written in their book.

6.2 Experiments

We analyzed 34 images of scenes, covering 9 different interpretations. In the sequel, we report the facts and rules used to interpret the pastoral scene, and the results of the performed tests; to make the paper more compact, for three more scenes, we only provide a textual explanation of the scene interpretation and the computed results. The Prolog rules for these three scenes can be found in the Appendix, along with five more examples. For sake of clarity, the `bb(X, Y, W, H)` argument of the `image` predicate is omitted in the following tables, which report the selected images and the respective interpretations with the test results.

6.3 Pastoral scene (corniforms group)

Interpretation of the scene by archeologists: A group of corniforms close to each others represents a pastoral scene.

Association between sub-image classification and sub-image interpretation:

```
interpretation('Corniform_Class', 'Corniform').
```

Rules for scene interpretation:

```
rule('Group_Of_Corniforms', Scenes) :-
    findall(scene(ID, BB, Cl, 'Corniform', Conf, SS),
           scene(ID, BB, Cl, 'Corniform', Conf, SS),
           Corns),
    sublist(Corns, Scenes),
    findall(BB,
           member(scene(_, BB, _, _, _, _), Scenes),
           BBs),
    prolog_list_to_java_list(BBs, JavaBBs),
    relations(GR),
    jpl_call(GR, group, [JavaBBs, 0.5], @(true)).
```

Explanation: The rule

- creates the set of corniforms in the scene by calling `findall(scene(ID, BB, Cl, 'Corniform', Conf, SS), scene(ID, BB, Cl, 'Corniform', Conf, SS), Corns)`,
- nondeterministically picks one partition of the set of corniforms by calling `sublist(Corns, Scenes)`,
- for the selected partition, retrieves the list of bounding boxes of the images therein by calling `findall(BB, member(scene(-, BB, -, -, -, -), Scenes), BBs)`,
- transforms the Prolog list BBs into a format suitable for being passed as an argument to a Java call (`prolog_list_to_java_list(BBs, JavaBBs)`), and finally
- checks if the bounding boxes form a group by calling `relations(GR), jpl_call(GR, group, [JavaBBs, 0.5], @(true))`.

Table 12 reports the results of the four analyzed images, all correctly interpreted.

6.4 Ritual sacrifice

Interpretation of the scene by archeologists: One halberd near one, or few more, corniforms, represents a ritual sacrifice. From the analysis of the available images, we identified three patterns: one where the BB of the corniform is inside the one of the halberd, another one where the two BBs are overlapping, and a last one where there are more corniforms.

Explanation: The rule shown in Section A.1 selects one halberd and another scene called Victim (a corniform or a group of corniforms) in the Scenes list. The check

Table 12. Results of the interpretation of Pastoral scenes

Image	Input Single Images	Resulting Interpretation	Final result
	<pre>image(0, [class('Corniform_Class', 1.0)]) image(1, [class('Corniform_Class', 1.0)]) image(2, [class('Corniform_Class', 1.0)]) image(3, [class('Corniform_Class', 1.0)]) image(4, [class('Corniform_Class', 1.0)]) image(5, [class('Corniform_Class', 1.0)]) image(6, [class('Corniform_Class', 1.0)]) image(7, [class('Solid_Ellipse_Class', 1.0)]) image(8, [class('Corniform_Class', 1.0)]) image(9, [class('Corniform_Class', 1.0)]) image(10, [class('Solid_Ellipse_Class', 1.0)]) image(11, [class('Corniform_Class', 0.6)])</pre>	<pre>I1 = [Group_Of_Corniforms(Corniform-0, Corniform-1, Corniform-2, Corniform-3, Corniform-4, Corniform-5, Corniform-6, Corniform-8, Corniform-9, Corniform-11). Cup_Stone_7, Cup_Stone_10] I2 = [Corniform-0, Corniform-1, Corniform-2, ..., Cup_Stone_7, Cup_Stone_10]</pre>	Passed
BL, R. 20, P. 5			
	<pre>image(0, [class('Corniform_Class', 1.0)]) image(1, [class('Corniform_Class', 1.0)]) image(2, [class('Corniform_Class', 1.0)]) image(3, [class('Corniform_Class', 1.0)]) image(4, [class('Corniform_Class', 1.0)])</pre>	<pre>I1 = [Group_Of_Corniforms(Corniform-0, Corniform-1, Corniform-2, Corniform-3, Corniform-4)] I2 = [Corniform-0, Corniform-1, Corniform-2, Corniform-3, Corniform-4]</pre>	Passed
BL, R. 20, P. 63			
	<pre>image(0, [class('Corniform_Class', 1.0)]) image(1, [class('Corniform_Class', 1.0)]) image(2, [class('Corniform_Class', 1.0)]) image(3, [class('Corniform_Class', 1.0)]) image(4, [class('Corniform_Class', 0.7)]) image(5, [class('Corniform_Class', 1.0)])</pre>	<pre>I1 = [Group_Of_Corniforms(Corniform-0, Corniform-1, Corniform-2, Corniform-3, Corniform-4, Corniform-5)] I2 = [Corniform-0, Corniform-1, Corniform-2, ...]</pre>	Passed
BL, R. 20, P. 80			
	<pre>image(0, [class('Corniform_Class', 1.0)]) image(1, [class('Corniform_Class', 1.0)]) image(2, [class('Corniform_Class', 1.0)]) image(3, [class('Corniform_Class', 1.0)])</pre>	<pre>I1 = [Group_Of_Corniforms(Corniform-0, Corniform-1, Corniform-2, Corniform-3)] I2 = [Corniform-0, Corniform-1, Corniform-2, Corniform-3]</pre>	Passed
DL, P. 230, F. 202			



Fig. 18. The High Goddess, DE, P. 328, F. 342.

succeeds if the halberd's BB contains or overlaps with the one of the Victim.

Table 13 reports the results of the four analyzed images: the last one has not been recognized because the two bounding boxes are neither overlapping nor one inside the other, as required by the rule.

6.5 Bull God birth

Interpretation of the scene by archeologists: One corniform below the High Goddess, shown in Figure 18, represents the Bull God born by the High Goddess. By analyzing the available images, two patterns were discovered: the first is where the High Goddess is above the Bull God, and close to him; the other is where she is above and partially overlaps with him.

Explanation: The rule shown in Section A.2 checks whether the token recognized as High Goddess is vertically aligned with the token representing the Bull God, and either overlaps with it, or it is close to it.

Table 14 reports the results of the four analyzed images: the fourth one has not been correctly interpreted because the High Goddess is not close enough to the Bull God. The problem might be easily solved by changing the proximity parameter in `jpl_call(GR, near, [BB1, BB2, 0.5], @(true))` from 0.5 to a higher value. Nevertheless, given that in most scenes representing the Bull God birth, the High Goddess is very close to him, increasing the proximity threshold might cause scenes with the High Goddess and one unrelated corniform nearby to be interpreted in the wrong way.

6.6 Storm God

Interpretation of the scene by archeologists: One dagger and one reticulum with some overlaps represent the Storm God.

Explanation: The rule shown in Section A.3 searches for a dagger and a reticulum, checking if they overlap.

Table 15 reports the results of four analyzed images; the first three ones have been correctly interpreted. The last one has not, because the reticulate and the dagger are very close, but do not overlap.

Table 13. Results of the interpretation of Ritual Sacrifice scenes



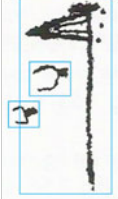

Image	Input Single Images	Resulting Interpretation	Final result
 DE, P. 187, F. 154(3)	<pre>image(0, [('Halberd_Class', 1.0)]) image(1, [('Corniform_Class', 1.0)])</pre>	<pre>I1=[Ritual_Sacrifice(Halberd-0,Corn-1)] I2=[Halberd-0, Corn-1]</pre>	Passed
 DE, P. 187, F. 154(10)	<pre>image(0, [('Halberd_Class', 1.0)]) image(1, [('Corniform_Class', 1.0)])</pre>	<pre>I1=[Ritual_Sacrifice(Halberd-0,Corn-1)] I2=[Halberd-0, Corn-1]</pre>	Passed
 DE, P. 187, F. 154(12)	<pre>image(0, [('Halberd_Class', 1.0)]) image(1, [('Corniform_Class', 1.0)]) image(2, [('Corniform_Class', 1.0)])</pre>	<pre>I1=[Ritual_Sacrifice(Halberd-0, Group_Of_Corniforms(Corniform-1,Corniform-2))] I2=[Halberd-0, Group_Of_Corniforms(Corn-1, Corn-2)] I3=[Halberd-0, Corn-1, Corn-2]</pre>	Passed
 BL, R. 8, P. 165	<pre>image(0, [('Halberd_Class', 1.0)]) image(1, [('Corniform_Class', 1.0)]) image(2, [('Corniform_Class', 1.0)])</pre>	<pre>I1=[Halberd-0, Group_Of_Corniforms(Corn-1, Corn-2)] I2=[Halberd-0, Corn-1, Corn-2]</pre>	Failed (no overlap)

Table 14. Results of the interpretation of Bull God Birth scenes





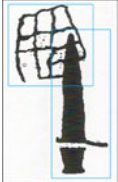
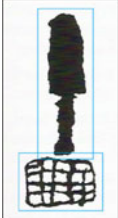

Image	Input Single Images	Resulting Interpretation	Final result
	image (0, [class('Double_Appendixes', 1.0)]) image (1, [class('Corniform_Class', 1.0)])	I1 = [HG_Giving_Birth_BG(High_Goddess-0, Bull_God-1)] I2 = [High_Goddess-0, Bull_God-1]	Passed
BL, R. 20, P. 134 	image (0, [class('Double_Appendixes', 1.0)]) image (1, [class('Corniform_Class', 1.0)])	I1 = [HG_Giving_Birth_BG(High_Goddess-0, Bull_God-1)] I2 = [High_Goddess-0, Bull_God-1]	Passed
DE, P. 330, F. 345(2) 	image (0, [class('Double_Appendixes', 1.0)]) image (1, [class('Corniform_Class', 1.0)])	I1 = [HG_Giving_Birth_BG(High_Goddess-0, Bull_God-1)] I2 = [High_Goddess-0, Bull_God-1]	Passed
DE, P. 330, F. 345(4) 	image (0, [class('Double_Appendixes', 1.0)]) image (1, [class('Corniform_Class', 1.0)])	I1 = [High_Goddess-0, Bull_God-1]	Failed (not close enough)
DE, P. 330, F. 345(3)			

Table 15. Results of the interpretation of Storm God scenes

Image	Input Single Images	Resulting Interpretation	Final result
 <p>DE, P. 171, F. 133(2)</p>	<pre>image(0, [('Dagger_Class', 1.0)]) image(1, [('Reticulum_Class', 1.0)])</pre>	<pre>I1=[Storm_God(Dagger-0, Reticulum-1)] I2=[Dagger-0, Reticulum-1]</pre>	Passed
 <p>DE, P. 171, F. 133(4)</p>	<pre>image(0, [('Dagger_Class', 1.0)]) image(1, [('Reticulum_Class', 1.0)])</pre>	<pre>I1=[Storm_God(Dagger-0, Reticulum-1)] I2=[Dagger-0, Reticulum-1]</pre>	Passed
 <p>DE, P. 171, F. 133(5)</p>	<pre>image(0, [('Dagger_Class', 1.0)]) image(1, [('Reticulum_Class', 1.0)])</pre>	<pre>I1=[Storm_God(Dagger-0, Reticulum-1)] I2=[Dagger-0, Reticulum-1]</pre>	Passed
 <p>BL, R. 20, P. 33</p>	<pre>image(0, [('Corniform_Class', 1.0)]) image(1, [('Reticulum_Class', 1.0)]) image(2, [('Dagger_Class', 1.0)]) image(3, [('Corniform_Class', 1.0)])</pre>	<pre>I1=[Dagger-2, Reticulum-1, Corn-0, Corn-3]</pre>	Failed (no overlap)

6.7 Discussion

Suitability of Prolog for modeling and implementing the scene interpretation rules. The power of Prolog for specifying scene interpretation rules is properly exemplified by the rule in Section 6.3 that exploits the `findall` all-solutions predicate for collecting all the images interpreted as corniforms into one set, generates one partition of the set in a nondeterministic way, and tests whether this partition enjoys the definition of being a group. If it does not, another partition is generated in backtracking and tested. By putting the `sublist` predicate inside a `findall` one, and then running the “is a group?” test on all the computed solutions, we would have obtained many more interpretations, one for each subgroup of corniforms in the scene. To keep it as simple and efficient as possible, the rule (`'Group_Of_Corniforms', Scenes`) goal succeeds as soon as the first group is found. While this rule was directly implemented in Prolog by the authors, based on the trivial intuition of what is a group of corniforms, other rules were sketched by the domain expert Dr. Nicoletta Bianchi using the formalism presented in Table 1, and then translated by the authors into Prolog, following translation rules that can be easily automatized. This is the case, for example, of the Bull God birth rule presented in Section A.2, whose rule in the user-friendly syntax is

```
rule('HG_Giving_Birth_BG', [X,Y]) {
  High_Goddess(X);
  Bull_God(Y);
  (vertical(X,Y) or near(X,Y) or overlap(X,Y))
}
```

Test results. We consider one test passed when OntoScene returns the correct interpretation, possibly together with other ones; 29 scenes out of 34 were correctly interpreted. The five scenes whose interpretation failed, did not satisfy the geometric constraints that the associated rule imposed. Failures are due to sub-images in the scene which do not overlap, while they should according to the rule, or that are not close enough, or that do not respect the expected orientation. In one case, failure is due to the lack of a suitable implementation for a geometric relation, “around”. Given that scenes in this domain present a high variability, even when they have been resorted to the same interpretation by the domain experts, writing “the perfect rules”, keeping them as compact as possible, and as few as possible, is very hard. For example, the last test presented in Section A.8 fails because the priest is above the repository, whereas the rule designed by the expert only accepts scenes where the priest is (or the two priests are) below. Adding one rule for coping with the failed test would not be difficult, but Nicoletta Bianchi knew that scenes like the one that failed the test are definitely less frequent than those that passed the test, and she suggested that – in some cases – obtaining a false negative could be better than designing many complex rules. In fact, OntoScene is meant to be a support to the domain experts, and not to substitute them in any way. Having a sound tool as OntoScene allows the expert to trust the “Passed” result and to check only the “Failed” one. Although the human in the loop is still required, this approach may save a lot of time.

Likelihood of interpretations. `SceneInterpreter` computes all the scene interpretations which are consistent with the provided rules, but says nothing on the likelihood of

one interpretation versus another. Coping with this further refinement does not represent a technical obstacle as it just resorts to sorting the elements in the list of computed interpretations according to some criterion. The actual obstacle is eliciting the sorting criterion from the domain experts and formalizing it. In all the 29 passed tests, the first interpretation returned, namely the one which “aggregates more” (see Section 5.4), turned out to be the correct one. This observation might suggest some heuristics for pruning the search tree, such as keeping the weight of the best interpretation obtained so far, and avoiding to expand branches whose weight is expected to be lower. However, the fact that this simple sorting criterion worked finely in the rock art application domain tells nothing on its generality. Different domain experts may have different personal opinions on how to select the correct interpretation of a scene, among many plausible ones, and associating a likelihood weight with each scene is not only domain-dependent, but even domain expert-dependent. This makes general and universally accepted sorting criteria difficult to assess: we did not face this issue in this paper, but it could be addressed either by integrating a heuristic criterion in the Algorithm X presented in Section 5.3 to stop recursion before the matrix M is empty, or by adding a post-processing stage of the `SceneInterpreter` output into the framework data flow. In the first case, the solution could be computed more efficiently, but could even get lost if the heuristic is not precise enough. In the second case, all the solutions should be computed, and efficiency would not benefit from the post-processing.

Performance. We did not assess the performance of `OntoScene`, both because efficiency was not our main concern, and because our experiments were run on scenes with no more than 11 sub-images: too few to raise efficiency issues. Despite the implemented optimization of Donald Knuth’s Algorithm X, where selection of the column to remove is made in a clever way, the complexity of the problem itself is high, and the only way to reduce it would be to give up finding the exact solution, and integrate some heuristics in the algorithm.

If stress-tested on scenes consisting of a large number of sub-images, we expect that `OntoScene` bottleneck should turn out to be `SceneInterpreter`, which would be a bottleneck even if implemented in any other language, because of the complexity of the exact cover algorithm it implements. [Dovier et al. \(2005\)](#) show how different NP-complete problems could be solved with either ASP ([Lifschitz 1999](#)) or CLP(FD) ([Marriott and Stuckey 1998](#)), and also on inputs with size greater than 2000. Based on these results, and considering that they date back to 15 years ago, we may suppose that, with today’s computing power, with efficient Prolog implementations, and possibly with a careful exploitation of more advanced technologies like ASP and CLP(FD), we could use `SceneInterpreter` on scenes with 2000 sub-images or more.

We point out, however, that adopting `OntoScene` to model scenes with hundreds or thousands of sub-images does not seem a viable approach to scene interpretation, and not because of performance issues. Rule modeling is worth the effort if the modeled rules are general enough to cover a large number of different scenes, but the more the scene elements, the more specific the rule. For example, designing an `OntoScene` rule for interpreting the scene represented in the Parthenon frieze would require to model the relations holding between/among 378 human figures and 245 animals. A precise rule for achieving this goal would succeed on the Parthenon frieze and would fail on anything else, and its usefulness would be very limited.

7 Conclusions and future works

OntoScene is a modular platform aimed at supporting the interpretation of complex scenes based on ontologies and logical rules defined in Prolog. Ontologies allow the designer to formalize the domain and make the system modular and interoperable with existing MASs, while Prolog provides a solid basis to define complex rules of interpretation in a way that can be affordable even for people with no background in Computational Logics. The feedback we got from Nicoletta Bianchi, with whom we designed the rules presented in Sections 6 and Appendix A, is that such rules are in a one-to-one, straightforward correspondence with the interpretation rules she had in mind, making their formalization easy to address at least in the user-friendly syntax presented in Section 3.

The overall design of our framework allows to easily change both the domain, modifying the ontology in the domain-specific parts (under *Classification* and *Interpretation* classes), the used geometric relationships, and the Prolog rules (that are formalized in an external file): furthermore, its inclusion in an already existing JADE MAS is quite simple (as described in Section 3.4) thanks to the adoption of the standard JADE usage of OWL ontologies. This makes the exploitation of our framework for other visual languages and existing systems easily achievable.

The case study presented in Section 6 comes from the IndianaMAS project. The results obtained from the experiments are encouraging and demonstrate the flexibility of our approach. The failures that we have reported might have been solved by minor changes to the rules or to the parameters therein. Given that the purpose of our experiments was neither to stress-test the framework, nor to provide a systematic evaluation of its precision and recall, but to show its applicability to a real domain, we left them as hints for a practical use of the framework.

Many improvements can be made to OntoScene.

So far, we assume that the Detector associates one bounding box with each sub-image: we did not take the possibility of detection ambiguity into account, as we assume that the Detector operates in a deterministic way. Apart from a growing time complexity, there would be no technical obstacles in allowing the Detector to produce more solutions (we mean producing, for the same input image, different decompositions into the sub-images detected there, namely different “sets of recognized bounding boxes”) and then deal with each of them separately, by running the Classifier and the SceneInterpreter on each of them.

Also, scenes are sensitive to orientation. While this is the correct approach in the rock art domain, where the interpretation may change depending, for example, on one sub-image being above or below another, it might turn out to be a limitation in other domains.

As far as Prolog rules are concerned, we only used rules meeting a very specific pattern: the initial part of the rule deals with the selection of the scenes to be aggregated, while the second part computes the geometric relations holding among them. This pattern worked well in the rock art domain, but more properties could be associated with images, ranging from features intrinsic to the image itself like the color, to semantically or emotionally related notions like the mood, and these properties could be part of the rules as well. OntoScene allows to add new properties to the *Image* class in the ontology and use these properties within the logical rules, according to the needs of the end user. For example,

we might want to extend the example presented in Figure 7 and define a happy, red warrior. The Prolog rule might be

```
rule('Red_Happy_Warrior', Scenes) :-
    scene(ID1, Img1, Class1, 'Human', Conf1, SS1),
    scene(ID2, Img2, Class2, 'Sword', Conf2, SS2),
    append([scene(ID1, Img1, Class1, 'Human', Conf1, SS1)],
           [scene(ID2, Img2, Class2, 'Sword', Conf2, SS2)], Scenes),
    bb(Img1, BB1),
    bb(Img2, BB2),
    mood(Img1, 'Happy'),
    color(Img2, 'Red'),
    relations(GR),
    jpl_call(GR, overlap, [BB1, BB2], @(true)).
```

where `mood` and `color` appear before the `relations` predicate.

Another extension we could address in the close future is to improve geometric relationships. OntoScene supports the addition and definition of new arbitrarily complex geometric relationships: the `Image` class in the ontology can be extended with new geometric properties as the area, the notion of `BB` can be refined using a polygonal closed line instead of a rectangle, and so on: the framework puts no limits on the type of accepted geometric relationships.

Finally, engraved rock art scenes are represented by black-white, bidimensional images often containing just a few elements placed in relatively simple geometric relationships. Given that the two phases of the `SceneInterpreter` computation (the creation of the scene graph and the generation of interpretations) are computationally heavy, they might require optimizations to scale to more complex domains. The possibility to improve the `SceneInterpreter` efficiency by rewriting it in ASP or CLP(FD) is under evaluation, although, before facing this language shift, we should find a domain where scenes are as complex as to motivate it.

The Prolog code for the `SceneInterpreter` and for some of the examples used for our experiments, and the OWL representation of the ontology, are currently available “as they are” from <http://www.disi.unige.it/person/MascardiV/Download/OntoScene.zip>. Once the above improvements will be ready, we plan to make OntoScene available to the research community via a well-designed website, after a suitable addition of comments, tutorials, and a user guide in English.

References

- AGUSTÍ, J., PUIGSEGUR, J. AND ROBERTSON, D. 1998. A visual syntax for logic and logic programming. *Journal of Visual Languages & Computing* 9, 4, 399–428.
- ANTANAS, L., VAN OTTERLO, M., MOGROVEJO, O., ANTONIO, J., TUYTELAARS, T. AND DE RAEDT, L. 2012. A relational distance-based framework for hierarchical image understanding. In *Proceedings of the 1st International Conference on Pattern Recognition Applications and Methods*, vol. 2, 206–218.
- BAADER, F., BÜRKERT, H.-J., HEINSOHN, J. AND HOLLUNDER, B. 1991. Terminological knowledge representation: A proposal for a terminological logic. In *International Workshop on Terminological Logics*. KIT-Report 89, TU Berlin, Fachbereich Informatik.

- BALDASSANO, C. 2015. *Visual Scene Perception in the Human Brain: Connections to Memory, Categorization, and Social Cognition*. Ph.D. thesis, Stanford University.
- BANNOUR, H. AND HUDELLOT, C. 2011. Towards ontologies for image interpretation and annotation. In *2011 9th International Workshop on Content-Based Multimedia Indexing (CBMI)*. IEEE, 211–216.
- BELLIFEMINE, F. L., CAIRE, G. AND GREENWOOD, D. 2007. *Developing Multi-Agent Systems with JADE*. John Wiley & Sons.
- BIANCHI, N. 2011. Mount Bego prehistoric rock carvings. *Adoranten*, 70–80.
- BICKNELL, C. 1913. *A Guide to the Prehistoric Rock Engravings in the Italian Maritime Alps*. Tip. G. Bessone.
- BRIOLA, D. 2016. *Agents and Ontologies for a Smart Management of Heterogeneous Data: The Indianamas System*. Studies in Computational Intelligence, vol. 616. Springer, 25–36.
- BRIOLA, D., DEUFEMIA, V., MASCARDI, V. AND PAOLINO, L. 2017. Agent-oriented and ontology-driven digital libraries: The Indianamas experience. *Software – Practice and Experience* 47, 11, 1773–1799.
- BRIOLA, D., DEUFEMIA, V., MASCARDI, V., PAOLINO, L. AND BIANCHI, N. 2014. Ontology-driven processing and management of digital rock art objects in Indianamas. In *Proceedings of 5th International Conference Digital Heritage (EuroMed 2014)*. Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), vol. 8740, 217–227.
- BRIOLA, D., MASCARDI, V. AND GIOSEFFI, M. 2018. OntologyBeanGenerator 5.0: Extending ontology concepts with methods and exceptions. In *Proceedings of the 19th Workshop From Objects to Agents, Palermo, Italy, M. Cossentino, L. Sabatucci and V. Seidita, Eds. CEUR Workshop Proceedings*, vol. 2215. CEUR-WS.org, 116–123.
- COSTAGLIOLA, G., DEUFEMIA, V. AND RISI, M. 2005. Sketch grammars: a formalism for describing and recognizing diagrammatic sketch languages. In *Eighth International Conference on Document Analysis and Recognition (ICDAR'05)*, vol. 2, 1226–1230.
- CRIMI, C., GUERCIO, A., NOTA, G., PACINI, G., TORTORA, G. AND TUCCI, M. 1991. Relation grammars and their application to multi-dimensional languages. *Journal of Visual Languages & Computing* 2, 4, 333–346.
- 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, 5, 969–1004.
- DE LUMLEY, H. AND ECHASSOUX, A. 2011. *La montagne sacrée du Bego*. CNRS Editions.
- DI MARTINO, B. AND ESPOSITO, A. 2016. A rule-based procedure for automatic recognition of design patterns in UML diagrams. *Software: Practice and Experience* 46, 7, 983–1007.
- DONAHUE, J., JIA, Y., VINYALS, O., HOFFMAN, J., ZHANG, N., TZENG, E. AND DARRELL, T. 2014. Decaf: A deep convolutional activation feature for generic visual recognition. In *International Conference on Machine Learning*, 647–655.
- DOVIER, A., FORMISANO, A. AND PONTELLI, E. 2005. A comparison of CLP(FD) and ASP solutions to np-complete problems. In *Logic Programming, 21st International Conference, ICLP 2005, Sitges, Spain, October 2–5, 2005, Proceedings*, M. Gabbriellini and G. Gupta, Eds. Lecture Notes in Computer Science, vol. 3668. Springer, 67–82.
- FORESTIER, G., DERIVAUX, S., WEMMERT, C. AND GAŃCARSKI, P. 2008. An evolutionary approach for ontology driven image interpretation. In *Workshops on Applications of Evolutionary Computation*. Springer, 295–304.
- GERBER, C., SIEKMANN, J. H., AND VIERKE, G. 1999. Holonic multi-agent systems. Tech. Rep. DFKI-RR-99-03, Deutsches Forschungszentrum für Künstliche Intelligenz - GmbH, Postfach 20 80, 67608 Kaiserslautern, FRG.
- GUARINO, N., OBERLE, D. AND STAAB, S. 2009. What is an ontology? In *Handbook on Ontologies*. Springer, 1–17.

- GUÉRIN, C., RIGAUD, C., BERTET, K. AND REVEL, A. 2017. An ontology-based framework for the automated analysis and interpretation of comic books' images. *Information Sciences* 378, C (Feb.), 109–130.
- HAARSLEV, V. 1999. A logic-based formalism for reasoning about visual representations. *Journal of Visual Languages and Computing* 10, 4, 421–445.
- HAARSLEV, V., MÖLLER, R. AND SCHRÖDER, C. 1994. Combining spatial and terminological reasoning. In *Annual Conference on Artificial Intelligence*. Springer, 142–153.
- HAARSLEV, V., MÖLLER, R. AND WESSEL, M. 2002. Visual spatial query languages: A semantics using description logic. In *Diagrammatic Representation and Reasoning*. Springer, 387–403.
- HAMMOND, T. AND DAVIS, R. 2007. LADDER, a sketching language for user interface developers. In *ACM SIGGRAPH 2007 Courses, SIGGRAPH '07*. ACM, New York, NY, USA.
- HE, K., ZHANG, X., REN, S. AND SUN, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- HELM, R. AND MARRIOTT, K. 1991. A declarative specification and semantics for visual languages. *Journal of Visual Languages & Computing* 2, 4, 311–331.
- HENDERSON, J. M. AND HOLLINGWORTH, A. 1999. High-level scene perception. *Annual Review of Psychology* 50, 1, 243–271.
- HILL, E. F. 2003. *Jess in Action: Java Rule-Based Systems*. Manning Publications Co.
- KARP, R. M. 1972. Reducibility among combinatorial problems. In *Proceedings of a Symposium on the Complexity of Computer Computations, held March 20-22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, USA*, R. E. Miller and J. W. Thatcher, Eds. The IBM Research Symposia Series. Plenum Press, New York, 85–103.
- KARTTUNEN, L. 1986. D-PATR: A development environment for unification-based grammars. In *Proceedings of the 11th Conference on Computational Linguistics*. Association for Computational Linguistics, 74–80.
- KNUTH, D. 2000. Dancing links. *Millennial Perspectives in Computer Science* 1.
- KONDO, H. M., VAN LOON, A. M., KAWAHARA, J.-I. AND MOORE, B. C. J. 2017. Auditory and visual scene analysis: an overview. *Philosophical Transactions* 372, 1714 (February).
- KVERAGA, K. AND BAR, M. Eds. 2014. *Scene Vision: Making Sense of What We See*. MIT Press.
- LADRET, D. AND RUEHER, M. 1991. Vlp: A visual logic programming language. *Journal of Visual Languages & Computing* 2, 2, 163–188.
- LI, S. AND YING, M. 2003. Region connection calculus: Its models and composition table. *Artif. Intell.* 145, 1–2, 121–146.
- LIFSCHITZ, V. 1999. Answer set planning. In *International Conference on Logic Programming and Nonmonotonic Reasoning*. Springer, 373–374.
- MARR, D. 1982. *Vision*. W.H. Freeman, San Francisco, CA.
- MARRIOTT, K. AND STUCKEY, P. 1998. *Programming with Constraints – An Introduction*. MIT Press.
- MASCARDI, V., BRIOLA, D., LOCORO, A., GRIGNANI, D., DEUFEMIA, V., PAOLINO, L., BIANCHI, N., DE LUMLEY, H., MALAFRONTI, D. AND RICCIARELLI, A. 2014. A holonic multi-agent system for sketch, image and text interpretation in the rock art domain. *Int. J. of Innovative Computing, Information and Control* 10, 1 (Feb.), 81–100.
- MCGUINNESS, D. L., VAN HARMELEN, F., *et al.* 2004. Owl web ontology language overview. *W3C Recommendation* 10, 10, 2004.
- MEYER, B. 1992. Pictures depicting pictures on the specification of visual languages by visual grammars. In *1992 IEEE Workshop on Visual Languages, 1992. Proceedings*. IEEE, 41–47.
- RANDELL, D. A. AND COHN, A. G. 1989. Modelling topological and metrical properties in physical processes. In *Proceedings of the 1st International Conference on Principles of Knowledge*

- Representation and Reasoning (KR '89)*. Toronto, Canada, May 15–18 1989., R. J. Brachman, H. J. Levesque, and R. Reiter, Eds. Morgan Kaufmann, 357–368.
- SANTOSH, K. C., LAMIROY, B., AND ROPERS, J. 2009. Inductive logic programming for symbol recognition. In *2009 10th International Conference on Document Analysis and Recognition*, 1330–1334.
- SIKOS, L. F. 2017. *Description Logics in Multimedia Reasoning*. Springer.
- SIMONYAN, K. AND ZISSERMAN, A. 2014. Very deep convolutional networks for large-scale image recognition. *CoRR abs/1409.1556*.
- ULLMAN, S. 1996. *High-Level Vision: Object Recognition and Visual Cognition*. MIT Press, Cambridge, MA.
- WAN, J., WANG, D., HOI, S. C. H., WU, P., ZHU, J., ZHANG, Y. AND LI, J. 2014. Deep learning for content-based image retrieval: A comprehensive study. In *Proceedings of the 22nd ACM International Conference on Multimedia*. ACM, 157–166.
- WITTENBURG, K., WEITZMAN, L. AND TALLEY, J. 1991. Unification-based grammars and tabular parsing for graphical languages. *Journal of Visual Languages & Computing* 2, 4, 347–370.
- WOOLDRIDGE, M. AND JENNINGS, N. R. 1995. Intelligent agents: theory and practice. *Knowledge Eng. Review* 10, 2, 115–152.

Appendix A

A.1 Ritual sacrifice

Association between sub-image classification and sub-image interpretation:

```
interpretation('Corniform_Class', 'Corniform').
interpretation('Halberd_Class', 'Halberd').
```

Rules for scene interpretation:

```
rule('Group_Of_Corniforms', Scenes):-...
    % Rule for interpreting a group of corniforms, Section 5.3.

rule('Ritual_Sacrifice', Scenes) :-
    scene(ID1, BB1, Class1, 'Halberd', Conf1, SS1),
    (Victim = 'Corniform';
     Victim = 'Group_Of_Corniforms'),
    scene(ID2, BB2, Class2, Victim, Conf2, SS2),
    append([scene(ID1, BB1, Class1, 'Halberd', Conf1, SS1)],
           [scene(ID2, BB2, Class2, Victim, Conf2, SS2)]), Scenes),
    relations(GR),
    (jpl_call(GR, contains, [BB1, BB2], @(true));
     jpl_call(GR, overlap, [BB1, BB2], @(true))).
```

A.2 Bull God birth

Association between sub-image classification and sub-image interpretation:

```
interpretation('Double_Appendixes', 'High_Goddess').
interpretation('Corniform_Class', 'Bull_God').
```

Rules for scene interpretation:

```
rule('HG_Giving_Birth_BG', Scenes) :-
    scene(ID1, BB1, Class1, 'High_Goddess', Conf1, SS1),
    scene(ID2, BB2, Class2, 'Bull_God', Conf2, SS2),
    append([scene(ID1, BB1, Class1, 'High_Goddess', Conf1, SS1)],
           [scene(ID2, BB2, Class2, 'Bull_God', Conf2, SS2)]),
    Scenes),
    relations(GR),
    jpl_call(GR, vertical, [BB1, BB2, 'up'], @(true)),
    (jpl_call(GR, near, [BB1, BB2, 0.5], @(true));
     jpl_call(GR, overlap, [BB1, BB2], @(true))).
```

A.3 Storm God

Association between sub-image classification and sub-image interpretation:

```
interpretation('Dagger_Class', 'Dagger').
interpretation('Reticulum_Class', 'Reticulum').
```

Rules for scene interpretation:

```
rule('Storm_God', Scenes) :-
  scene(ID1, BB1, Class1, 'Dagger', Conf1, SS1),
  scene(ID2, BB2, Class2, 'Reticulum', Conf2, SS2),
  append([scene(ID1, BB1, Class1, 'Dagger', Conf1, SS1)],
         [scene(ID2, BB2, Class2, 'Reticulum', Conf2, SS2)]), Scenes),
  relations(GeometricRelations),
  jpl_call(GeometricRelations, overlap, [BB1, BB2], @(true)).
```

A.4 Rain invocation

Interpretation of the scene by archeologists: One human wielding a halberd or an axe (or in general a weapon) represents the rain invocation. From the analysis of the available images, we discovered that the weapon is usually on above the human, positioned in a vertical or diagonal way.

Association between sub-image classification and sub-image interpretation:

```
interpretation('Human_Class', 'Human').
interpretation('Halberd_Class', 'Halberd').
interpretation('Axe_Class', 'Axe').
```

Rules for scene interpretation:

```
rule('Rain_Summon', Scenes) :-
  scene(ID1, BB1, Class1, 'Human', Conf1, SS1),
  subclass_of('Weapon_Class', Class),
  interpretation(Class, Weapon),
  scene(ID2, BB2, Class2, Weapon, Conf2, SS2),
  append([scene(ID1, BB1, Class1, 'Human', Conf1, SS1)],
         [scene(ID2, BB2, Class2, Weapon, Conf2, SS2)]), Scenes),
  relations(GR),
  jpl_call(GR, vertical, [BB2, BB1, 'up'], @(true)),
  !,
  jpl_call(GR, near, [BB1, BB2, 0.5], @(true)).

rule('Rain_Summon', Scenes) :-
  [omissis] % as in the previous rule
  relations(GR),
  (jpl_call(GR, diagonal, [BB2, BB1, 'ne'], @(true)) ;
   jpl_call(GR, diagonal, [BB2, BB1, 'nw'], @(true))),
  jpl_call(GR, near, [BB1, BB2, 0.5], @(true)).
```





Explanation: The rule searches for a human figure, then it searches for a weapon (note that halberd and axe are subclasses of weapon in the domain ontology, so we write a general rule including all the weapons as required by the archeologists) and checks for the correct geometrical relationship; then, the rule checks if the BBs of the human and of the weapon are close to each other, and if the one of the weapon is above the human, in vertical or diagonal relationship.

Table A1 reports the results of the four analyzed images, all correctly interpreted.

A.5 Queens fight

Interpretation of the scene by archeologists: Two corniforms with juxtaposed horns represent a ritual fighting called in archeology “the Queens Fight”. The two

Table A 1. *Results of the interpretation of Rain Invocation scenes*

Image	Input Single Images	Resulting Interpretation	Final result
 <p data-bbox="365 505 517 523">DE, P. 189, F. 157(4)</p>	<p data-bbox="553 352 877 371">image(0, [('Human_Class', 1.0)])</p> <p data-bbox="553 373 887 392">image(1, [('Halberd_Class', 1.0)])</p>	<p data-bbox="966 352 1333 371">I1=[Rain_Summon(Human-0, Halberd-1)]</p> <p data-bbox="966 373 1202 392">I2=[Human-0, Halberd-1]</p>	Passed
 <p data-bbox="365 742 517 760">DE, P. 189, F. 157(6)</p>	<p data-bbox="553 622 877 640">image(0, [('Human_Class', 1.0)])</p> <p data-bbox="553 643 887 661">image(1, [('Halberd_Class', 1.0)])</p>	<p data-bbox="966 622 1333 640">I1=[Rain_Summon(Human-0, Halberd-1)]</p> <p data-bbox="966 643 1202 661">I2=[Human-0, Halberd-1]</p>	Passed
 <p data-bbox="365 873 517 892">DE, P. 189, F. 157(12)</p>	<p data-bbox="553 806 877 825">image(0, [('Human_Class', 1.0)])</p> <p data-bbox="553 827 887 846">image(1, [('Halberd_Class', 1.0)])</p>	<p data-bbox="966 806 1333 825">I1=[Rain_Summon(Human-0, Halberd-1)]</p> <p data-bbox="966 827 1202 846">I2=[Human-0, Halberd-1]</p>	Passed
 <p data-bbox="365 1127 517 1146">DE, P. 200, F. 171(2)</p>	<p data-bbox="553 997 877 1015">image(0, [('Human_Class', 1.0)])</p> <p data-bbox="553 1018 846 1037">image(1, [('Axe_Class', 1.0)])</p>	<p data-bbox="966 997 1290 1015">I1=[Rain_Summon(Human-0, Axe-1)]</p> <p data-bbox="966 1018 1164 1037">I2=[Human-0, Axe-1]</p>	Passed

corniforms must be one over the other, with contrary directions of the horns (we assume that the Classifier is able to discriminate between the two different positions), and their BBs may, or may not, intersect, but should be close to each other.

Association between sub-image classification and sub-image interpretation:

```
interpretation('Up_Corn_Class', 'Corniform').
interpretation('Up_Down_Corn_Class', 'Corniform').
```

Rules for scene interpretation:

```
rule('Queens_Fight', Scenes) :-
  scene(ID1, BB1, 'Up_Down_Corn_Class', 'Corniform', Conf1, SS1),
  scene(ID2, BB2, 'Up_Corn_Class', 'Corniform', Conf2, SS2),
  append([scene(ID1, BB1, 'Up_Down_Corn_Class', 'Corniform', Conf1, SS1)],
         [scene(ID2, BB2, 'Up_Corn_Class', 'Corniform', Conf2, SS2)], Scenes),
  relations(GR),
  jpl_call(GR, vertical, [BB1, BB2, 'up'], @(true)),
  jpl_call(GR, near, [BB1, BB2, 0.5], @(true)).
```

Explanation: The rule searches for two corniforms, one with up horns and the other with down horns, in vertical relationship and close to each other.

Table A 2 reports the results of the four analyzed images: the last one is not correctly interpreted because the geometrical relationships “Around” has not been implemented yet, and a third unexpected element (a rock) appears in the scene.

A.6 Bull God

Interpretation of the scene by archeologists: One corniform inside the horns of another one represents the Bull God. By analyzing the available images, two patterns were discovered: the first is one or more corniforms inside another one, and another is a group of corniforms vertically aligned, not necessary one inside the other.

Association between sub-image classification and sub-image interpretation:

```
interpretation('Up_Corn_Class', 'Corniform').
```

Rules for scene interpretation:

```
rule('Group_Of_Corniforms', Scenes) :- .... (from Section 5.3)

rule('Bull_God', Scenes) :- (Inner = 'Corniform';
  Inner = 'Group_Of_Corniforms'),
  scene(ID1, BB1, Class1, Inner, Conf1, SS1),
  scene(ID2, BB2, Class2, 'Corniform', Conf2, SS2),
  append([scene(ID1, BB1, Class1, Inner, Conf1, SS1)],
         [scene(ID2, BB2, Class2, 'Corniform', Conf2, SS2)], Scenes),
  relations(GR),
  jpl_call(GR, contains, [BB2, BB1], @(true)).

rule('Bull_God', Scenes) :-
  get_corniforms_same_direction(Corniforms),
  sublist(Corniforms, Scenes, length(Scenes, Len), Len > 1,
  findall(BB, member(scene(_, BB, _, _, _), Scenes), BBs),
  relations(GR),
  test_vertical(BBs, GR).
```

Explanation: The first rule selects one corniform (or a group of corniforms) and another one from the list and checks if they are one inside the other. The second rule uses the predicate `get_corniforms_same_direction` to get all the corniforms with the same orientation and checks if they are in vertical relationships. We omit here the definition of the `test_vertical(BBs, GR)` predicate.

Table A 2. *Results of the interpretation of Queens Figh*

Image	Input Single Images	Resulting Interpretation	Final result
 <p>DE, P. 221, F. 191(5)</p>	<p>image (0, [('Up_Down_Corn_Class', 1.0)])</p> <p>image (1, [('Up_Corn_Class', 1.0)])</p>	<p>I1=[Queens_Fight (Corniform-0, Corniform-1)]</p> <p>I2=[Corniform-0, Corniform-1]</p>	Passed
 <p>DE, P. 221, F. 191(7)</p>	<p>image (0, [('Up_Down_Corn_Class', 1.0)])</p> <p>image (1, [('Up_Corn_Class', 1.0)])</p>	<p>I1=[Queens_Fight (Corniform-0, Corniform-1)]</p> <p>I2=[Corniform-0, Corniform-1]</p>	Passed
 <p>DE, P. 221, F. 191(8)</p>	<p>image (0, [('Up_Down_Corn_Class', 1.0)])</p> <p>image (1, [('Up_Corn_Class', 1.0)])</p>	<p>I1=[Queens_Fight (Corniform-0, Corniform-1)]</p> <p>I2=[Corniform-0, Corniform-1]</p>	Passed
 <p>DE, P. 221, F. 191(9)</p>	<p>image (0, [('Up_Down_Corn_Class', 1.0)])</p> <p>image (1, [('Up_Corn_Class', 1.0)])</p> <p>image (2, [('Rock_Class', 1.0)])</p>	<p>I1=[Corniform-0, Corniform-1, Rock-1]</p>	Failed (Around not avail- able)

Table A3 reports the results of the four analyzed images, all correctly interpreted.

A.7 Rain propitiatory rite

Interpretation of the scene by archeologists: One dagger between the horns of a corniform represents a propitiatory rite for the rain. The two sub-images should intersect and at the same time the dagger should be partially inside the horns, above them. With the currently implemented geometrical relationships we cannot express this relation in a precise way, so we approximated it.

Association between sub-image classification and sub-image interpretation:

```
interpretation('Corniform_Class', 'Corniform').
interpretation('Dagger_Class', 'Dagger').
```

Rules for scene interpretation:

```
rule('Rain_Propitiatory_Rite', Scenes) :-
  scene(ID1, BB1, Class1, 'Dagger', Conf1, SS1),
  scene(ID2, BB2, Class2, 'Corniform', Conf2, SS2),
  append([scene(ID1, BB1, Class1, 'Dagger', Conf1, SS1)],
         [scene(ID2, BB2, Class2, 'Corniform', Conf2, SS2)], Scenes),
  relations(GR),
  jpl_call(GR, vertical, [BB1, BB2, 'up'], @(true)),
  jpl_call(GR, overlap, [BB1, BB2], @(true)).
```

Explanation: The rule searches for a dagger and a corniform, checking if they overlap and if the dagger is above the corniform.

Table A4 reports the results of the three analyzed images, all correctly interpreted.

A.8 Agricultural rite

Interpretation of the scene by archeologists: One or two priests making water spring from an artificial repository represent an agricultural rite. The most recurring pattern includes one or two humans holding a repository, which is above them, from which the water falls down.

Association between sub-image classification and sub-image interpretation:

```
interpretation('Human_Class', 'Priest').
interpretation('Repository_Class', 'Repository').
interpretation('Water_Class', 'Water').
```

Rules for scene interpretation:

```
rule('Agricultural_Rite',
     [scene(ID1, BB1, Class1, 'Priest', Conf1, SS1),
      scene(ID2, BB2, Class2, 'Priest', Conf2, SS2),
      scene(ID3, BB3, Class3, 'Repository', Conf3, SS3),
      scene(ID4, BB4, Class4, 'Water', Conf4, SS4)]) :-
  scene(ID1, BB1, Class1, 'Priest', Conf1, SS1),
  scene(ID2, BB2, Class2, 'Priest', Conf2, SS2),
  scene(ID3, BB3, Class3, 'Repository', Conf3, SS3),
  scene(ID4, BB4, Class4, 'Water', Conf4, SS4),
  relations(GR),
  jpl_call(GR, diagonal, [BB1, BB3, 'sw'], @(true)),
  jpl_call(GR, near, [BB1, BB3, 0.5], @(true)),
  jpl_call(GR, diagonal, [BB2, BB3, 'se'], @(true)),
  jpl_call(GR, near, [BB2, BB3, 0.5], @(true)),
  jpl_call(GR, vertical, [BB3, BB4, 'up'], @(true)),
  jpl_call(GR, near, [BB3, BB4, 0.5], @(true)).
```


Table A 3. *Results of the interpretation of Bull God*



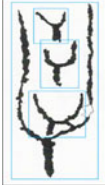
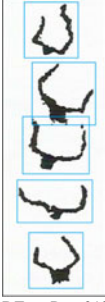
Image	Input Single Images	Resulting Interpretation	Final result
 <p>DE, P. 218, F. 186(2)</p>	<pre>image(0, [('Up_Corn_Class', 1.0)]) image(1, [('Up_Corn_Class', 1.0)])</pre>	<pre>I1=[Bull_God(Corn-0,Corn-1)] I2=[Corn-0,Corn-1]</pre>	Passed
 <p>DE, P. 218, F. 186(9)</p>	<pre>image(0, [('Up_Corn_Class', 1.0)]) image(1, [('Up_Corn_Class', 1.0)]) image(2, [('Up_Corn_Class', 1.0)])</pre>	<pre>I1=[Bull_God(Corn-0,Corn-1,Corn-2)] I2=[Corn-0,Corn-1,Corn-2]</pre>	Passed
 <p>DE, P. 218, F. 186(7)</p>	<pre>image(0, [('Up_Corn_Class', 1.0)]) image(1, [('Up_Corn_Class', 1.0)]) image(2, [('Up_Corn_Class', 1.0)]) image(3, [('Up_Corn_Class', 1.0)])</pre>	<pre>I1=[Bull_God(Corn-0,Corn-1,Corn-2,Corn-3)] I2=[Corn-0, Bull_God(Corn-1,Corn-2,Corn-3)] I3=[Corn-0,Corn-3,Bull_God(Corn-1, Corn-2)] I4=[Corn-0,Corn-1,Corn-2,Corn-3]</pre>	Passed
 <p>DE, P. 218, F. 186(18)</p>	<pre>image(0, [('Up_Corn_Class', 1.0)]) image(1, [('Up_Corn_Class', 1.0)]) image(2, [('Up_Corn_Class', 1.0)]) image(3, [('Up_Corn_Class', 1.0)]) image(4, [('Up_Corn_Class', 1.0)])</pre>	<pre>I1=[Bull_God(Corn-0,Corn-1,Corn-2,Corn-3, Corn-4)] I2=[Corn-0,Bull_God(Corn-1, Corn-2,Corn-3,Corn-4)] I3=[Bull_God(Corn-0, Corn-1,Corn-2), Bull_God(Corn-3, Corn-4)] ... other 7 interpretations</pre>	Passed

Table A4. Results of the interpretation of Rain Propitiatory Rite

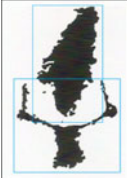
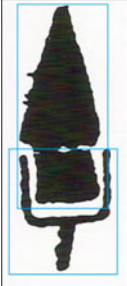




Image	Input Single Images	Resulting Interpretation	Final result
 <p>DE, P. 170, F. 132(1)</p>	<pre>image(0, [('Dagger_Class', 1.0)]) image(1, [('Corniform_Class', 1.0)])</pre>	<pre>I1=[Rain_Propitiatory_Rite(Dagger-0, Corniform-1)] I2=[Dagger-0, Corniform-1]</pre>	Passed
 <p>DE, P. 170, F. 132(2)</p>	<pre>image(0, [('Dagger_Class', 1.0)]) image(1, [('Corniform_Class', 1.0)])</pre>	<pre>I1=[Rain_Propitiatory_Rite(Dagger-0, Corniform-1)] I2=[Dagger-0, Corniform-1]</pre>	Passed
 <p>DE, P. 170, F. 132(7)</p>	<pre>image(0, [('Dagger_Class', 1.0)]) image(1, [('Corniform_Class', 1.0)])</pre>	<pre>I1=[Rain_Propitiatory_Rite(Dagger-0, Corniform-1)] I2=[Dagger-0, Corniform-1]</pre>	Passed

Table A 5. Results of the interpretation of Agricultural Rite

Image	Input Single Images	Resulting Interpretation	Final result
 <p>DE, P. 284, F. 268(1)</p>	<p>image (0, [('Human_Class', 1.0)]) image (1, [('Human_Class', 1.0)]) image (2, [('Repository_Class', 1.0)]) image (3, [('Water_Class', 1.0)])</p>	<p>I1=[Agricultural_Rite(Priest-0, Priest-1, Repository-2, Water-3)] I2=[Priest-0, Priest-1, Repository-2, Water-3]</p>	Passed
 <p>DE, P. 284, F. 268(4)</p>	<p>image (0, [('Human_Class', 1.0)]) image (1, [('Repository_Class', 1.0)]) image (2, [('Water_Class', 1.0)])</p>	<p>I1=[Agricultural_Rite(Priest-0, Repository-1, Water-2)] I2=[Priest-0, Repository-1, Water-2]</p>	Passed
 <p>DE, P. 284, F. 268(5)</p>	<p>image (0, [('Human_Class', 1.0)]) image (1, [('Repository_Class', 1.0)]) image (2, [('Water_Class', 1.0)])</p>	<p>I1=[Human-0, Repository-1, Water-2]</p>	Failed (Human over repository)

Explanation: The rule searches for the two humans, the water and the repository, checking if the two humans are in diagonal (one on the left and one on the right) below the repository, and if the water is under the repository. All the images should be close to each other. Another rule, searching for only one human, is not reported since it is very similar to one shown here.

Table A 5 reports the results of the three analyzed images: the last one has not been correctly interpreted because the human is above (not below and in diagonal) the repository.