Cognitive Modelling of Event Ordering Reasoning in Imagistic Domains[1]

Laura S. Mastella[1], Mara Abel[1], Luís C. Lamb[1], Luiz F. De Ros[2]
[1]Instituto de Informática and [2]Instituto de Geociências
Universidade Federal do Rio Grande do Sul
Porto Alegre, RS 91501-970, Brazil
{mastella,marabel,lamb,lfderos}@inf.ufrgs.br

Abstract

The inference of temporal information from past event occurrences in imagistic domains is relevant in several applications in knowledge engineering. In such applications, the order in which events have happened is imprinted in the domain as visual-spatial relations among its elements. Therefore, the interpretation of the relative ordering in which those events have occurred is essential for understanding the domain evolution. We propose a cognitive model for event ordering reasoning within domains whose elements have been modified by past events. From the analysis of cognitive abilities of experts we propose new ontology constructs for knowledge modelling associated to Problem-Solving Methods. We illustrate the effectiveness of the model by means of an application to an imagistic domain.

1 Introduction

The inference of temporal information from past event occurrences in imagistic domains is relevant in several applications in knowledge engineering [Thagard and Shelley, 1997]. There are several approaches to model temporal reasoning in intelligent systems such as algebras, logics, calculi and effective mechanisms to reason about time [Allen, 1983; Gabbay *et al.*, 1994]. Nonetheless, the development of ontological languages providing adequate support to formalise event-based temporal inferences in such domains is still incipient. In addition, most proposals consider an absolute notion of time, which can be of limited use in a number of applications.

However, here we aim at deriving relative temporal information from another dimension (the visual-spatial relations between the elements of the domain), which is foreign to most proposals developed so far. We are interested in proposing a cognitive model of the reasoning of an expert about an imagistic domain, when deriving a temporal relation from a visual-spatial relation.

Imagistic domains are characterised by visual-spatial relations requiring from the problem-solver the ability of applying visual recognition of objects, and from this initial recognition, to start the search and analytical methods in order to interpret these objects. Such abilities define a form of *imagistic reasoning* [Yip and Zhao, 1996], which cannot de described only in terms of the geometrical properties of objects, but by the meaning and significance that this objects have in selecting a particular reasoning path in a visual-spatial domain.

In some domains, the recognition of key features of elements and the identification of the visual-spatial relations among them is an important research issue [Ericsson and Smith, 1991]. For instance, a geologist identifies visual-spatial relations among rock constituents (called paragenetic relations), as does a physician when analysing medical images to identify a pathology.

In order to build an ontology and appropriate problem-solving methods for such domains, a long process of knowledge acquisition from experts was carried out. The analysis of the cognitive abilities of the expert led to the development of new ontology constructs for knowledge modelling associated to Problem-Solving Methods (PSM). These constructs and methods are then shown to be capable of modelling the expert's reasoning when deriving the sequence of events which led to the visual-spatial organisation of the domain under analysis.

Section 2 presents the basics of ontologies and their extensions with temporal constructs and a brief description of the PSMs that shall be used. Section 3 defines the main constructs and primitives of the cognitive model for event reasoning and PSMs to infer sequences of events. Section 4 describes the application of the developed models to an imagistic domain, namely *sedimentary petrography*. Section 5

[1] This work is supported by the Brazilian Research Council CNPq through the CTPetro programme grant 502009/2003-9. L.C. Lamb is partly supported by CAPES and CNPq.

presents an initial validation of the model. Section 6 concludes the paper and discusses directions for future research.

2 Preliminaries

An ontology is a formal explicit specification of a shared conceptualisation [Studer *et al.*, 1998]. Usually, domain ontologies represent static or declarative knowledge about a domain: the main concepts, their attributes, the relationships between them, axioms, rules, etc [Gómez-Pérez *et al.*, 2004]. [Gruber, 1993] has identified five types of ontological constructs: *classes*, which represent elements of a domain defined by a set of attributes and their possible values; *relations*, which represent the types of association between domain elements; *functions*, special relations that map one or several elements of the domain to a unique element; *axioms*, which are statements that are always true about the domain; and *instances*, the individuals of the ontology.

While ontologies describe the structural and static part of knowledge, the dynamic part is described using a PSM [Gómez-Pérez and Benjamins, 1999]. A PSM describes the reasoning process of a knowledge-based system (KBS), specifying the knowledge and data required by an inference process at a more abstract and structured level. This reasoning pattern is modelled by three related parts: (i) a *competence specification* related to the solution of a task, (ii) an *operational specification* described by high-level modelling primitives and (iii) *requirements/assumptions* of the method in terms of domain knowledge.

Recently, some authors have aimed at augmenting the expressive power of ontologies representing temporal information. A causal time ontology is proposed by [Kitamura *et al.*, 1997], in which an event represents instantaneous changes of qualitative values of parameters and their resulting values at a time point. The work of [Zhou and Fikes, 2000] treats both time points and time intervals as primitive elements of a time line, and the classes, relations, axioms and instances of the ontology are built upon those primitives. The extension of DAML to DAML-Time aims to develop an ontology of time that expresses temporal aspects common to any formalisation of time, such as temporal relations on instants and intervals [Hobbs and Pustejovsky, 2003]. [Bennett and Galton, 2004] propose a highly expressive language for representing temporal relationships and events, called VEL (Versatile Event Logic). A single VEL formula can contain both time-point and interval variables. In most works, the addition of an explicit temporal dimension to ontology constructs is typically defined by an absolute time stamp associated with objects of the ontology, i.e., time points or time intervals.

However, in several application domains, events are not to be interpreted as having time stamps labelling them. It is the *relative ordering* in which events have occurred that is of fundamental importance in understanding, analysing, and inferring the evolution of elements in such domains. The aim of our approach is on representing how the domain evolves through a sequence of events that act as operators transforming the state-space of a domain.

The cognitive model proposed here shall be applied for modelling domains that present evolving elements which have been modified by events that have occurred in an unplanned order. The order in which the events happened is imprinted in the domain as complex visual-spatial relations among the elements of the domain. The cognitive models we introduce here intend to represent how to identify the events from the domain elements' characteristics, and how to find the order in which the events have occurred out of the visual-spatial relations in such domains.

We deal with visual-spatial knowledge translated into a symbolic representation, instead of using typical numerical approaches for image processing. This follows a line of research outlined in [Kosslyn, 1994] which describes that some features (and types) of visual-spatial knowledge are dealt by the human brain as symbolic entities, not image representations; therefore we believe that these features can be represented by means of a symbolic approach. Moreover, the proposed cognitive models are not intended to recognise the elements of the domain, but to infer new information from their symbolic description made by the user of the expert system.

3 A Cognitive Model for Event Ordering

Let us now define the required constructs to deal with event ordering and the associated problem-solving methods.

3.1 An Ontology of Events

We propose an extension to the constructs of ontological representation - classes, relations, inference rules, axioms and instances - for evolving domains in order to capture the meaning of *events* and *temporal relations between them*. Such proposed constructs shall be applied for modelling domains that need to infer past sequences of occurrences to provide a better comprehension of the current state of the domain. Further, we aim at understanding how the domain elements were produced and by which events, and also at formalising the explanation through a representation of an event sequence. We define the new constructs as follows:

- *Events* are class-transforming constructs. They represent phenomena that generate or modify the elements of the domain. They are characterised by specific domain-dependent attributes, but not by a time stamp. Events are also described by rules that associate them to their products. Events are associated to each other by an ordering relation.
- *Temporal relation*: a construct proposed to represent the ordering relation between events. We have defined the binary relations *before, after* and *during* in order to reflect the ordering between events.

Furthermore, in our domain ontology, we extend the notion of *inference rules* to make them more expressive. They express functions between instances of relations, in addition to functions between instances of classes.

Most knowledge representation structures used in domain ontologies define rules as expressions about the *attributes* of objects, as in the following rule:

if *classA.attribute1 = value-x* ***then*** *classB.attribute2 = value-y*

We use this type of rule to model functions that identify events from their produced elements. The characteristics of the elements (expressed by class attributes in the ontology) are used to indicate the event that originated or modified the element. The instances of these rules define that when a particular value for an attribute is found in the description of an element then a determined event is indicated as the generating event. These inference rules are called *event indication rules*. In addition, we define an extended type of inference rule in knowledge representation ontologies. The purpose is to express functions between instances of *relations*, instead of classes. This second type of rule, referred to as *temporal implication rule*, is defined to allow the inference of binary *temporal relations* between events from *visual-spatial relations* between the elements, as in

if *visual_relation(A,B)* ***then*** *temporal_relation(A,B)*

Next, we present a PSM to model inferences about the ordering of events in ontologies built with these new types of constructs and inference rules.

3.2 A PSM for Ordering Events

In order to show that the proposed constructs are adequate and sufficient for inferring event ordering, the reasoning method was modelled as a PSM, abstracted from the heuristics used by a highly-trained expert in the domain. The PSM uses the temporal constructs to identify and order the events that originated the current characteristics of the domain.

The *competence* of our PSM takes a description of domain elements as input, identifies the events that acted over these elements and infers the ordering in which such events have occurred. The *assumptions* of our PSM are: (1) the characteristics imprinted in the domain elements indicate the event that generated the element; (2) the way elements are spatially disposed reflects the order in which they were originated in the domain. The *requirements* are the inference rules that associate characteristics of domain elements to their generating events and the inference rules that associate visual-spatial and ordering relations. The input of the PSM is a description of the visual-spatial features without any temporal reference. The sequence of events is inferred as the PSM output. The *operational specification* describes the inferences. The inferences are applied to the elements in the domain knowledge (represented as knowledge roles) in order to derive new information. The inferences used in our PSM can be detailed as follows:

- *Abstract*: select attributes that describe the relevant aspects of the element for problem-solving.
- *Relation abstract*: takes an element as input and outputs relevant relations between such element and others.
- *Specify*: takes a relation as input and outputs the related element (the second argument of the relation).
- *Cover*: given a set of rules and objects that match the premise of the rules, outputs the conclusion of the rule.
- *Order*: takes a pair of events and the ordering relation between them and produces an ordered pair of events.

The *event ordering model* is shown in Figure 1, as a graphic model that presents the knowledge roles, the information flow, and the inferences. It is explained as follows.

The knowledge role (label 1) represents elements in the application domain that is under analysis. The elements of the domain are related to each other through complex visual-spatial relations (label 2), which are abstracted from the element description. As visual-spatial relations are binary, the other argument of the relation can be specified as the related element (label 3), which is also a domain element. The result of these steps is a pair of domain elements associated by a visual-spatial relation.

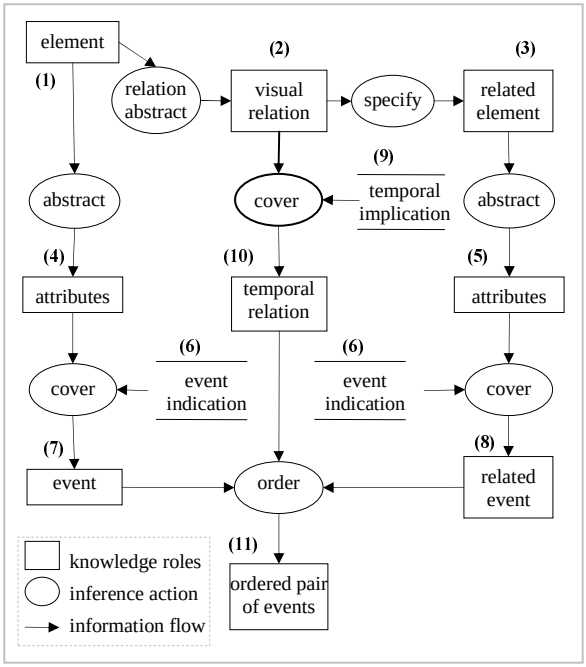

Figure 1: Problem-Solving Method for Event Ordering.

Every domain element is described by a set of attributes which defines its generating event. Relevant attributes (labelled as 4 and 5) are abstracted from the set of attributes. Event indication rules (label 6) cover the generating events (labelled as 7 and 8) from the attributes of the element.

The visual-spatial relation between domain elements is used to infer a temporal relation (label 10) using the temporal implication rules (label 9). From the temporal relation it is possible to order the events, producing an ordered pair of events (label 11). A detailed example of application of this reasoning method is presented in Section 5.

Thus, the above PSM can be used in domain ontologies modelled using the proposed temporal constructs with the objective of inferring temporal relations from visual-spatial relations in an imagistic domain.

4 Modelling an Imagistic Domain

We are now in a position to illustrate the application of the proposed models to an imagistic domain scenario.

4.1 Describing the Petrographic Domain

We consider *sedimentary petrography*, which aims at evaluating the economic prospects of oil fields and reservoirs by means of interpretation techniques from descriptions of rock samples. One particular technique widely used to determine the quality of a reservoir is called *diagenetic sequence interpretation* [Anjos *et al.*, 2000], which aims to infer the order in which diagenetic events occurred in a rock. *Diagenetic events* are physical-chemical processes that act over the sediments transforming them into solid rocks and modifying the porosity and permeability of a potential oil-reservoir. The visual-spatial relations among rock constituents reflect the changes undergone by the rock as a result of diagenetic events. In Figure 2 we show an example of rock sample and the visual-spatial relations between minerals (called *paragenetic relations*) that can be identified.

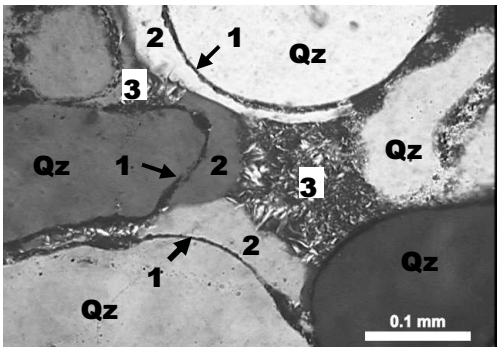

Figure 2: Detail view of a rock sample with (1) hematite covering quartz grains (Qz); (2) quartz covering the hematite; (3) quartz being covered by illite, which is covering the hematite.

Even though the arrangement is said to be spatial, they cannot be translated into simple spatial relations, since they are by themselves a result of a visual interpretation of particular rock aspects. This interpretation requires previous experience in mineral recognition that goes beyond the analysis of geometrical and topological characteristics. Figure 2 illustrates examples of paragenetic relations.[2]

The expert in petrography points out the ordering of events by observing how the constituents are spatially and visually related to each other, based on his extensive previous knowledge. Using a simple example, if one mineral appears to be on top of other mineral, it means that the former was generated in the rock later than the latter.

Some interpretations techniques used for the evaluation of oil reservoirs had already been modelled in the *PetroGrapher* system, an intelligent database application to support the description and interpretation of sedimentary rock samples [Abel *et al.*, 2004]. The vocabulary of petrography was elicited as a result of previous work on the domain [Abel *et al.*, 2003] and modelled as a domain ontology mostly structured as a partonomy[3] of classes. The system implements interpretation tasks as PSMs over the domain ontology.

Although the petrography domain ontology had already been modelled, temporal aspects of the domain needed in the interpretation of diagenetic sequences were still lacking. These aspects had to be elicited later, by means of knowledge acquisition (KA) sessions with an expert in sedimentary petrography. In order to acknowledge the heuristics of the expert when performing the diagenetic sequence interpretation, several knowledge elicitation techniques have been applied [Cooke, 1994]. Firstly, we applied interviews and observation techniques to elicit the key vocabulary used in the interpretation, to determine the task's purpose and to understand how the interpretation is realised. Then, in order to reveal the expert's implicit knowledge (i.e. knowledge used by the expert that is not easily verbalised [Polanyi, 1974]), we applied card sorting and repertory grids techniques. Thus, the main results obtained in the KA session were insights about temporal aspects that needed to be modelled (such as diagenetic events and ordering relations) and an abstraction of the expert's reasoning steps when solving the interpretation task.

Next, we shall model the task of diagenetic sequence interpretation (i.e. how to identify the order in which events have occurred) by using the constructs proposed in Section 3 above, so as to represent the temporal aspects identified in the domain. We shall also show how to instantiate the PSM for event ordering of Figure 1 to infer the sequence of diagenetic events that occurred in a rock.

4.2 From the Event Ontology to the Petrography Domain

Let us now show how the petrography domain ontology was extended with temporal aspects using the new temporal constructs.

a) Rock Constituents (modelled as *classes*). Rock constituents are minerals and pores that form a rock. They are represented in the ontology using the *class* construct, since they are elements of the domain. The characteristics of the constituents are described as *attributes* of the class that represents it. It is important to represent these features since they reflect how this constituent was produced or modified, indicating its generating diagenetic events.

b) Paragenetic Relations (modelled as *relations*). Paragenetic relations are visual-spatial arrangements among constituents. Since a paragenetic relation is defined between a pair of constituents, the construct used in the ontology is a *binary relation*.

c) Diagenetic Events (modelled as *events*). These events are physical-chemical phenomena of rock consolidation and are represented in the model as an *event*. The expert does

2 There are several kinds of paragenetic relations that can be identified such as *covering, replacing, dissolving*. Each particular paragenetic relation creates visual-spatial relations between different types of constituents. For instance, the relation *dissolution* occurs only between a pore and a diagenetic constituent.

3 Partonomy is a domain organization based on a part-of relation rather than on a kind-of (is-a) relation [Martínez-Béjar and Fernández-Breis, 2000]. E.g., a foot is part of a leg which is part of a body.

not take into account the absolute period of time the event has happened, but only the order in which the diagenetic event has occurred in relation to other events.

d) Ordering Relation (modelled as a *temporal relation*). A rock is formed as a result of several different diagenetic events, which can happen in a simultaneous or in a sequential way. To simplify the computational treatment of the sequence, we treat the ordering of events in pairs, as does the expert. The relations between pairs of events were represented using the *temporal relation* construct to model the temporal relations *after, before,* and *during.*

e) Inference rules (modelled as *functions*). As explained in Section 3.1, the expert is able to indicate the generating events by analysing the characteristics of the constituents. For instance, when the attribute *modifier* of a constituent holds the value *deformed*, it is possible to conclude that the event that transformed the constituent is *compaction*. Hence, it was necessary to represent this knowledge as *event indication rules*. These inference rules define an association between *constituents* and *diagenetic events*, e.g. (R1) below:

if *constituent.modifier = deformed*
then *event.event_name = compaction* R1

After identifying the events, the expert is able to infer the order in which they occurred by means of the visual-spatial (paragenetic) relations among the constituents. For instance, when a mineral appears to be *covering* (on top of) another mineral, the expert says that the event that formed the first mineral occurred later than the event that formed the latter. The first part of this particular expert's rule is assuming a *paragenetic relation* between constituents. The second part is defining an *ordering relation* between events. Thus, we need to represent this knowledge as *temporal implication rules*, defined in Section 3.1. An example of this type of rule is the following (R2):

if *covering(constituent1, constituent2)*
and *produced_by(constituent1, event1)*
and *produced_by(constituent2, event2)* R2
then *after(event1, event2)*

If *constituent1* is covering *constituent2* and the events that produced them are, respectively, *event1* and *event2*, then the conclusion is that *event1* happened after *event2*.

5 Initial Validation of the Cognitive Models

The temporal constructs and the proposed PSM have been applied to the petrography domain. They were implemented as an inference module within the *PetroGrapher* system, called *diagenetic sequence interpretation module*. Real rock samples were described by the expert in the *PetroGrapher* system and he also provided a previous interpretation of the sequence of diagenetic events. Thus, the interpretation produced by the algorithm was compared to the expert's interpretation of the sequence. In order to illustrate this initial validation, a step-by-step interpretation of a sample is described in what follows.

As input to the algorithm we have a description of a rock sample, which is composed by 132 attributes that detail the description of the constituents of the rock. The algorithm analyses each constituent individually. For example, the description of one constituent is as follows:

[quartz - intergranular discontinuous pore-lining - covering <diagenetic-constituents> - hematite]

In the PSM depicted in Figure 1, this line corresponds to the knowledge role labelled as *element* (1). The next steps are: abstracting the paragenetic relation (*covering <diagenetic-constituents>*) from the description (role *visual relation* - label 2) and specifying the constituent visually related to the former (role *related element* - label 3). Thus, the related constituent is *hematite*, described as follows:

[*hematite – coating - intergranular continuous pore lining - within intergranular primary porosity - macroporosity intergranular*]

The result of these steps is a pair of constituents associated by a paragenetic relation (*quartz - covering - hematite*).

A rock constituent is characterised by some attributes which provide an indication of the generating event of the constituent. The algorithm abstracts such attributes from the description of the constituent. The attributes are the roles labelled as 4 and 5 in Figure 1. The relevant attributes are used as premises for the event indication rules (label 6), returning the event that produced or modified the constituent. These last inference steps are performed for both constituents, resulting in a pair of diagenetic events (labelled as 7 and 8). For the constituents considered in this example, the following events are inferred by the rules: *precipitation of quartz* and *precipitation of hematite.*

The paragenetic relation between constituents (label 2) is used as premise for the temporal inference rules (label 9). The result is the ordering relation (label 10) (e.g. given in rule R2, where the paragenetic relation *covering* implies the temporal relation *after*). The ordering relation and the pair of events are used as inputs for the inference named *order*, which results an ordered pair of events (label 11), which, in this case is

after(Cementation of Quartz, Cementation of Hematite)

The algorithm that mechanizes the PSM is implemented as a loop that stops when there are no more constituents to be evaluated. The algorithm outputs a list of ordered pairs, e.g. the following list for the sample being evaluated:

after(Cementation of Hematite, Deposition of Detrital Quartz)
after(Cementation of Quartz, Cementation of Hematite)
before(Cementation of Quartz, Cementation of Illite)

We use a directed graph to create a sequence of events from the list of ordered pairs. An edge from the event that comes before (in the source vertex) to the event that comes later in time (in the target vertex) defines the ordering relation. For instance, the result of mapping the list above to a directed graph is represented in Figure 3.

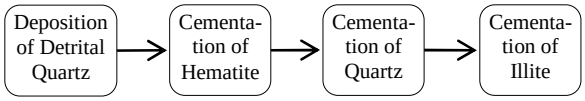

Figure 3: Events sequence produced by the inference algorithm

Such inferences depend on how clear the visual-spatial arrangements of constituents are in a particular sample.

However, in some cases, not even the expert is able to produce a complete sequence of events, because some paragenetic relations may be visible (and then described) in one sample, but not in another one. Thus, for some rock samples the algorithm produces sequences of events that are not totally connected leading to an incomplete interpretation of the sequence.

Although the resulting sequence may sometimes be incomplete, the produced outcome is certainly relevant to the domain. Any sequence of events that can be inferred from a rock description is essential in understanding how the porosity and the permeability of the rock were affected, and how this influences the quality of the oil reservoir.

6 Conclusion

We have presented a new cognitive model for reasoning about event ordering relations in imagistic domains. The proposed model is composed of an ontology of events and an associated Problem-Solving Method. From extensive sessions of knowledge acquisition from an expert we were able to: (i) identify the need for new temporal aspects in a domain ontology, which led us to the formalisation of such aspects as new ontological constructs to represent relative notions of time - so far lacking in the literature; (ii) to abstract the PSM that models a process of event ordering interpretation, and (iii) to identify the need of extending the notion of inference rules in an ontology so as to allow the inference of temporal relations from visual-spatial relations.

The proposed model aims to mimic the reasoning process of an expert who infers events that generated the current domain characteristics from complex visual-spatial relations between domain elements. We have illustrated the effectiveness of the model in a real-world application and initially validated it through the implementation of the model in an expert system. Further experiments leading towards a thorough validation of the system are currently being carried out. We believe that it is possible to use the proposed models in other domains, e.g. to infer the sequence of events that may have caused climate changes by analysing ice testimonies in glaciology. Further, in archaeology it is of fundamental importance to define a sequence of events in the history of the object that led to its current shape and structure [Thagard and Shelley, 1997].

Finally, this work can also be seen as a starting point towards investigating *abductive procedures* that correspond to the reasoning of the expert. This would corroborate Peirce's original claims regarding abduction as a form of visual reasoning [Thagard and Shelley, 1997]. In particular, the models presented here model the way an expert explains and interprets visual-spatial relations in search for the best explanation about the sequence of events that caused them.

References

[Abel *et al.*, 2003] M. Abel, J. Campbell and L. F. D. Ros. Knowledge acquisition techniques for visual expertise: A study of oil-reservoir evaluation. In *Proc. of the 2nd ICPI - IJCAI03*, pages 74-79, Acapulco, 2003.

[Abel *et al.*, 2004] M. Abel, L. A. L. Silva, L. F. D. Ros, L. Mastella, J. Campbell and T. Novello. Petrographer: Managing petrographic data and knowledge using an intelligent database application. *Expert Systems with Applications,* 26: 9-18, 2004.

[Allen, 1983] J. F. Allen. Mantaining knowledge about temporal intervals. *Communications of the ACM,* 26: 832-843, 1983.

[Anjos *et al.*, 2000] S. M. C. Anjos, L. F. De Ros, R. S. Souza, C. M. A. Silva and C. L. Sombra. Depositional and diagenetic controls on the reservoir quality of lower cretaceous sandstones. *Am. Ass. Petrl. Geol. Bull.,* 84: 1719-1742, 2000.

[Bennett and Galton, 2004] B. Bennett and A. Galton. A unifying semantics for time and events. *Artif. Intell.,* 153: 13-48, 2004.

[Cooke, 1994] N. Cooke. Varieties of knowledge elicitation techniques. *Int. Jnl. Hum.-Comp. Studies,* 41: 801-849, 1994.

[Ericsson and Smith, 1991] K. Ericsson and J. Smith. *Toward a general theory of expertise: Prospects and limits.* Cambridge U. Press, 1991.

[Gabbay *et al.*, 1994] D. M. Gabbay, I. Hodkinson and M. Reynolds. *Temporal logic*. Oxford U. Press, 1994.

[Gómez-Pérez and Benjamins, 1999] A. Gómez-Pérez and V. R. Benjamins. Overview of knowledge sharing and reuse components: Ontologies and problem-solving methods. In *Proc. of KRR-5 - IJCAI99*, pages 1.1-1.15, 1999.

[Gómez-Pérez *et al.*, 2004] A. Gómez-Pérez, M. Fernández-López and O. Corcho. *Ontological engineering*. Springer, 2004.

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

[Hobbs and Pustejovsky, 2003] J. Hobbs and J. Pustejovsky. Annotating and reasoning about time and events. In *Proc. AAAI Spr. Symp. Logical Form. Commonsense Reas.*, 2003.

[Kitamura *et al.*, 1997] Y. Kitamura, M. Ikeda and R. Mizoguchi. A causal time ontology for qualitative reasoning. In *Proc. of the 15th IJCAI*, pages 501-507, 1997.

[Kosslyn, 1994] S. M. Kosslyn. Resolving the imagery debates. *Image and brain: The resolution of the imagery debate*. S. M. Kosslyn (ed.), MIT Press**:** 2-23, 1994.

[Martínez-Béjar and Fernández-Breis, 2000] R. Martínez-Béjar and J. T. Fernández-Breis. A cooperative tool for facilitating knowledge management. *Expert Systems with Applications,* 18: 315-330, 2000.

[Polanyi, 1974] M. Polanyi. *Personal knowledge*. The Univ. of Chicago Press, 1974.

[Studer *et al.*, 1998] R. Studer, V. R. Benjamins and D. Fensel. Knowledge engineering: Principles and methods. *Data & Knowledge Engineering,* 25: 161-197, 1998.

[Thagard and Shelley, 1997] P. Thagard and C. Shelley. Abductive reasoning: Logic, visual thinking, and coherence. *Logic and scientific methods*. M. D. Chiara (ed.), Kluwer**:** 413-427, 1997.

[Yip and Zhao, 1996] K. Yip and F. Zhao. Spatial aggregation: Theory and applications. *Jnl. of Art. Intel. Research,* 5: 01-26, 1996.

[Zhou and Fikes, 2000] Q. Zhou and R. Fikes. *A reusable time ontology*, Knowledge Systems Lab, Stanford U., 2000.