Abel, M. Silva, L.A.L , Campbell, J.A. De Ros, L.F. *Knowledge acquisition and interpretation problem-solving methods for visual expertise: a study of petroleum-reservoir evaluation*. Journal of Petroleum Science and Engineering. To appear in Spring 2005.

Knowledge acquisition and interpretation problem-solving methods for visual expertise: a study of petroleum-reservoir evaluation

Mara Abel[a], Luís A. Lima Silva[a,b], John A. Campbell[b], Luiz F. De Ros[c]

[a]Instituto de Informática, Universidade Federal do Rio Grande do Sul; Av. Bento Gonçalves, 9500; Campus do Vale; Caixa Postal: 15064; CEP: 91501-970; Porto Alegre, RS, Brasil

[b]Department of Computer Science, University College London; Gower Street, London; WC1E 6BT, UK

[c]Instituto de Geociências, Universidade Federal do Rio Grande do Sul; Av. Bento Gonçalves, 9500; Campus do Vale; Caixa Postal: 15001; CEP: 91501-970; Porto Alegre, RS, Brasil

Abstract

The paper presents a practical study of expertise in sedimentary petrography and its incorporation into a knowledge-based computing system, *PetroGrapher*. The study identifies features of expert geologists' essentially visual knowledge that limit the usefulness of traditional logic-oriented symbolic methods for developing relevant knowledge-based systems. In particular, the experts' expressions of knowledge differ significantly from what is taught through the literature about their subject. Also, the means (ontology) of describing the explicit objects of geological data are not well suited to describing how experts combine such objects and make inferences about them. The paper indicates how to deal with those two different levels of ontology, proposes "knowledge graphs" as an effective medium for linking them, and explains their place in a general descriptive model of how experts solve problems in interpretation based on visual evidence. The model is being validated through the use of *PetroGrapher*, for management of data and knowledge about reservoir rocks, in a petroleum company.

Keywords: Knowledge engineering, knowledge acquisition, expertise, problem-solving method, sedimentary petrography

1 Introduction

The *PetroGrapher* project was started with the aim of producing a knowledge-based computing system which would embody both extensive data about sedimentary rocks and usable representations of expert knowledge in sedimentary petrography, e.g. to help novice geologists improve their competence in interpreting sedimentary rocks, and to make the management of significant volumes of petrographic information more effective. In addition, we intend the system to be used in a petroleum-company environment, to collate this information in a standard way, interpret it automatically or semi-automatically with the help of the embodied knowledge, combine it with

Authors: [a] Fax: +55 (51) 33167308, E-mail address: marabel@inf.ufrgs.br; [b] Fax: +44 (20) 73871397, E-mail address: {l.silva, j.campbell}@cs.ucl.ac.uk; [c] Fax: +55 (51) 33191811, E-mail address: lfderos@inf.ufrgs.br.

information from other sources, preserve it in the database of the organization, and share it through the organization's distributed computing facilities.

Although the design and implementation of the *PetroGrapher* knowledge-base system was not a simple exercise, the work of identifying and capturing the knowledge in a form amenable to use in a knowledge-base system has been an even more complicated experience. The present paper gives an account of the major parts of our overall work, with emphasis on the concepts and the lessons that are most likely to be of value to those who wish to study geological or other subject-areas where the expertise is essentially visual and to build knowledge-base systems to incorporate such expertise.

The special character of visually-based or image-based expertise is noted in section 2, with reference to sedimentary geology. Capturing this (or any) knowledge for knowledge-base system use should now be a systematic activity; section 3 outlines the issues of that activity: knowledge engineering. In section 4 we focus on petrographic knowledge, and discuss the outcomes of experimental observation of geologists' treatment of visual information. Section 5 deals with interpretation of sedimentary rock specimens in particular, taking advantage of the foundations laid in the previous sections. The nature of expertise in sedimentary petrography and its demands on the structure of a suitable knowledge-base system are discussed in section 6. An overview of the implementation of the *PetroGrapher* knowledge-base system is then given.

2 A view of geological knowledge

Geological interpretation involves the recognition of the products created by natural effects and the construction of a theory about how the natural forces have worked to generate these products. Geology is strongly based on visual interpretation of features imprinted in the rock or land by physical phenomena. Since this process of interpretation is improved mainly by practical experience, not by supervised learning typical of school classes, geology is the kind of domain in which experts develop themselves as strategic sources of knowledge.

Expertise is even more crucial in petroleum exploration, since expert knowledge can substantially decrease the risks of exploration, and increase the efficiency of hydrocarbon production. Geologists apply high levels of expertise to answer questions such as: What kinds of reservoir rocks should be sought? Which are the criteria and parameters to be used to evaluate the reservoir? Which methods are needed for petroleum field development? Capturing this knowledge and representing it in a formal model is a permanent aim for knowledge management in petroleum companies.

During the exploration for new petroleum deposits and the development and production of oilfields, two sets of information have fundamental importance. The first refers to the external geometry of the reservoir in describing the shape and extension of the sedimentary units that contain petroleum, and the second refers to the internal structure and characteristics of the reservoir rocks. Although not essentially, most of the data relevant for geological interpretation of petroleum reservoirs consist of visual information that cannot be described symbolically through geometric components, such as size and format, alone. Many of the aspects recognized by a geologist during the interpretation task have no formal denomination and are learnt through an implicit process during training and field experience.

The use of features without names in supporting problem solving is a current practice in many natural domains that demand pattern matching as a part of the interpretation problem-solving process. We call these objects visual chunks, following the work of De Groot in the chess domain (De Groot 1965). These objects constitute the implicit body of knowledge that refers to the unarticulated knowledge that someone applies in daily tasks, but that the user is not able to describe in words, according to Polanyi in (Polanyi 1966). Articulated or explicit knowledge refers to consciously recognized objects and how these objects are organized. This portion of knowledge in the context of Artificial Intelligence (AI) is called ontology. Implicit (which we also call *tacit*, an adjective that became popular after being used by Nonaka and Takeuchi (Nonaka, Takeuchi et al. 1995)) and explicit knowledge should be seen just as two separate aspects of knowledge and not essentially different manifestations of it, although it is evidently useful to treat them in separate ways, especially for knowledge acquisition.

The word *ontology* has been taken from the technical vocabulary of philosophy, where it means a systematic explanation of the world which is shared by a community. Nowadays, the best-known definition of *ontology* is that of Gruber, as slightly modified by Borst in (Borst 1997): “Ontologies are defined as a formal specification of a shared conceptualization”. This definition implies that an ontology is an abstract model of a part of the world, described by a vocabulary defined in some unambiguous way (such as in some computer-processed language). “Shared” means that an ontology captures consensual knowledge that is used by a community, and not just by an individual (Studer, Benjamins et al. 1998).

Extracting the ontology of a domain is one important objective of knowledge acquisition techniques; it is sometimes misunderstood to be the only one. Along with ontology, which represents the explicit part of knowledge, it is necessary to identify the implicit knowledge applied by experts, and also the dynamic mechanisms of problem-solving. In geology, as in many other natural domains, the role of implicit knowledge in the expert problem-solving process has a dominating influence on the results achieved. Therefore, it deserves special attention in knowledge engineering.

The main focus of the expert task in this application is the problem of rock interpretation, which is concerned with the seeking of a reasonable explanation for the formation process of a sedimentary rock. According to (Harmon and Sawyer 1990) *interpretation* means analysis of data to determine their meanings. When applied to visual domains, in keeping with Abel in (Abel 2001), image interpretation deals with the matching of more abstract and diagnostic visual features than geometric data. The geometric features are not used on their own to produce relevant inferences.

Petrographic interpretation is carried out by the analysis of the relationship of rock features that are discerned through inspection under an optical microscope. These features, collected by some visual pattern-matching process, are selected/combined and only afterwards named and organized for incorporation into an explicit body of knowledge. The common sequence of inference steps and the domain objects manipulated by inferences are together modeled as a problem-solving method (PSM) (Gómez-Pérez and Benjamins 1999).

A PSM is basically a template describing the steps of what an expert does in solving a problem - here, interpretation of visual evidence. It is particularized by filling out the specific details of the kind of visual evidence (rock samples) that interests us. The filling-out process, guided and made easier by the template, is central to systematic knowledge acquisition. The end result can then be used in implementation of a knowledge-based system, as it captures in a suitably structured way both the data and the knowledge required for geological inference.

Knowledge engineering, which is the activity that involves both acquisition and structuring of the knowledge in question, is outlined in the next section.

3 Knowledge engineering

The main objective of knowledge engineering is to transform the *ad hoc* process of building knowledge-based systems into a discipline of engineering, based on methods, languages and specialized tools (Studer, Benjamins et al. 1998). Following the idea of a Knowledge Level as proposed by Newell (Newell 1982) which was further reinterpreted in the context of knowledge-based system development by van de Velde in (Velde 1993), the knowledge acquisition process has started to be seen as a modeling process, in opposition to the traditional simplistic view that, to construct an expert system, knowledge need only be transferred directly from an expert to some computer.

Modern knowledge engineering methodologies such as PROTÉGÉ (Puerta, Egar et al. 1992) and CommonKADS (Schreiber, Akkermans et al. 2000) (which is the one applied in this project) are successful in providing structure for the development process of knowledge-based systems by identifying intermediate models and defining the languages and organization of these models. These approaches are still lacking, however, in provision of the same degree of methodological support for the knowledge acquisition process, especially when the main source is the human expert (when acquisition is then called elicitation).

Knowledge acquisition amounts to the process of collection, elicitation, interpretation and formalization of the data regarding the functioning of expertise in a particular domain. Its objective is to reduce the communication gap between the expert or knowledge worker and the knowledge engineer, allowing the knowledge to become independent of its sources. The traditional motivation for knowledge acquisition, which underlies the development of any knowledge-based system, is complemented by other objectives from software engineering – accurate capturing of user requirements - and knowledge management – collecting and sharing the information of knowledge workers within organizations.

The main classes of knowledge acquisition techniques are described briefly below.

- Interviews, observation and protocol analysis: A grouping of many different techniques that demand direct interaction with experts through retrospective or concurrent inquiry. In the retrospective interview, the expert narrates a memory of how a problem was solved. This description is commonly rationalized and omits many crucial details. In concurrent interviewing, e.g. via observation and protocol analysis, the expert verbalizes his/her reasoning *during* the problem-solving process while it is being recorded and observed. The result is more trustworthy than a direct interview, but the expert is usually unable to verbalize what he/she is doing when the inference requires both a high level of abstract reasoning and a low (i.e. concrete) level of sensory activity. The collected information is commonly imperfect and needs to be complemented through further techniques.

- Classification techniques aim to identify the terms and concepts of the domain and how these concepts are organized in classes, groups or components, according to the expert. These include card-sorting and multidimensional scaling techniques (Wright and Ayton 1987).

- Collecting cases: A general label for all techniques that exploit recorded cases in knowledge acquisition, such as scenario analysis, recovering of events, and the analysis of legacy cases for use in case-base reasoning systems (Leake 1996).

- Extracting cause-effect relations. This includes techniques used to extract causal relations among concepts of the domain (such as evidence for conclusions, or problems and their applicable solutions). Repertory grids (Kelly 1955), rule-extraction, knowledge graphs (Leão and Rocha 1990), (Abel, Castilho et al. 1998) and conceptual graphs (Sowa 1984) belong in this class.

- Identifying the reasoning pattern. Problem-solving methods (PSM) and inference structures (Schreiber, Akkermans et al. 2000) are graphical representations of the inference process involved in problem solving, described at an abstract (though not so abstract as to be generic) level.

The accepted knowledge acquisition techniques are effective in revealing the ontology underlined in expert reasoning, or the explicit component of knowledge. However, little progress on elucidating the unarticulated (tacit) parts of expert knowledge has occurred, e.g. visual recognition or the integration of perceived objects in the domain ontology. These issues are the main focus of our present work.

The growing importance of the modeling approach in modern knowledge engineering has put into circulation several distinct formalisms to express the declarative part of knowledge that is captured as the domain ontology. These formalisms replace the previously-used models drawn from classical AI (where they were never proposed because of any particular consideration of knowledge engineering), such as frames (Fikes and Kehler 1985), production rules (Buchanan and Shortliffe 1984), semantic networks (Reboh 1981), among others. The main paradigms of languages currently used to represent ontologies are Description Logic (Baader, Calvanese et al. 2003) and Frame Logic (Angele and Lausen 2004). These main paradigms are complemented by RDF (Resource Description Framework) (Brickley and Guha 2000) and its evolution, the OWL (Web Ontology Language) (Antoniou and Harmelen 2004) mainly applied in connection with the "Semantic Web" scheme. Knowledge engineering methodologies, such as PROTÉGÉ and CommonKADS, propose their own knowledge-typed models to describe declarative knowledge, mainly based on the descriptions of concepts, attributes, domains of attributes, relations and rules.

Our project has applied the CommonKADS formalism to represent knowledge, extending it to represent abstract visual objects and interpretation PSMs. The PSM describes the reasoning process of a knowledge-based system in an implementation-independent way, through use of the 3 components: i) a competence specification related to the solution of a task, ii) an operational specification described by high-level modeling primitives and iii) requirements/assumptions of the method in terms of domain knowledge. We have developed the rock interpretation PSM accordingly, which allows us to explain the way that knowledge and data are employed to solve problems of rock interpretation.

4 Investigation of expertise in petrography

In the present study, our intention has been to determine the essential details of the cognitive process and objects that support expertise in sedimentary petrography, by using tests which are similar to those found occasionally in cognitive psychology, e.g. in (Ericsson and Smith 1991). The study was designed to understand a duality of expert behavior when explaining the domain and arriving at decisions. When describing rocks or presenting the domain, expert petrographers are formal and use precise vocabulary, but on the other hand they are not able to give even a semi-formal justification of their interpretations based on their own petrographic descriptions. This behavior is a serious obstacle to the extraction of causal relationships in knowledge acquisition. Elucidating how experts organize their domain knowledge and which objects are applied during the solution of common problems, and comparing these with the corresponding behavior of novices, should provide useful insights into the nature of geological expertise and how to train aspirants to expert-level competence.

Our investigation was conducted over a group of 19 geologists with distinct levels of expertise in sedimentary petrography. The group was selected among lecturers, undergraduate and graduate students of the Geosciences Institute of Universidade Federal do Rio Grande do Sul, and geologists from a Brazilian petroleum company. Practical experience, in place of theoretical knowledge alone, was a fundamental prerequisite for the selected group. All of the selected geologists utilize petrography as a common working tool. The members of this group were first classified as novices, intermediates or experts. Novices were students or geologists who had received at least 100 hours of training in sedimentary petrography. Intermediates were geologists who employed petrography as a daily work tool. Experts possessed at least 10 years of experience in the subject and utilized sedimentary petrography directly for more than 10 hours per week. The analyzed sample was distributed, using these criteria, in a bar graphic from the highest to the lowest level of expertise and compared with the measured behaviors demonstrated in the experiment. The most significant of these graphs are presented in Figures 1 to 3. The whole experiment has been described thoroughly in (Abel 2001).

The group was requested to carry out 5 different tests, based on the presentation of images from rocks, using a high-definition video system attached to an optical microscope. The first experiment was designed to evaluate long-term memory. Participants were requested to describe fully a thin section that they had examined one hour previously. The second test requested also a full description of a thin section, but without time restriction, in order to evaluate the richness of technical vocabulary. The third test investigated short-term memory, by requiring description of a rock image just previously shown for a very short time (30 seconds). In the fourth experiment, the recall of a first set of pictures of common objects (animals, landscapes and people) and a second set of sedimentary rocks was requested just after these had been shown for a short time. In the fifth experiment, the geologists were asked to divide another set of pictures into subsets (i.e., classify them) and explain their criteria for this subdivision.

The set of experiments was conceived in order to measure the degrees of association of 3 indicators with the predefined class of expertise. The indicators were:

(1) amount of significant information in the description obtained after or during image exposition: it was expected that experts would know more about the domain and that their knowledge would be expressed through the use of technical and precise vocabulary;

(2) intensity of use of interpreted features, the *visual chunks,* instead of features having objective geometrical properties: it was expected that experts would develop image recognition at a more abstract level than novices and would be able to demonstrate this in recognizing features that needed only a short process of inference for identification. For example, we considered "dolomite cementation" as a geometric, plain feature, although "marine cementation" required some level of interpretation to be defined. Also, "massive dolomite intragranular displacive" as a geometric feature, but "calcrete" is an interpretation of the same feature;

(3) efficiency of organization and indexing of the domain: it was expected that experts would be more effective in grouping and classifying new information related to the domain and could demonstrate this ability in experiments involving their memory.

4.1 Description of results

In the first and second experiments, we asked geologists to describe a set of sandstone images collected by an optical microscope in two stages: while they were examining the image and well after they had examined it. The descriptions produced for the 6 rock samples were then considered together, and the vocabulary and significant vocabulary were counted for each of the geologists. We found no relation between the quantities of words or even of significant words in descriptions produced by experts and by novices. This is illustrated in Figure 1. This picture presents (along the vertical axis) the quantity of words and the quantity of significant words in the description from experts (E), intermediates (I) and novices (N), previously sorted according to the level of expertise along the horizontal axis. By significant words, we mean the applied terms after eliminating all prepositions, common verbs, adverbs and other expressions not related to the domain. On the average, experts record more efficiently the details of the rock inspected, but the overall relation is not so simple: some novices exhibit an expert level of performance and some experts perform like novices. As suggested by the experiment, knowledge can not be connected in a simple monotonic way with the size of vocabulary in a domain ontology.

When considering the use of visual chunks in the description, experts clearly demonstrate a higher average, as illustrated in Figure 2, which shows the amount of interpreted expressions applied by geologists in the same two experiments associated with Figure 1. Here, the vertical axis in Figure 2 indicates how many clearly interpreted features were used in describing the rock. The general pattern was the same in the subsequent short-term memory experiments involving sedimentary rock samples. In those experiments, where the geologists were required to described the rock after a short time of display of the image (about 30 seconds), the faster the expert is requested to classify or interpret rocks, the more he/she will try to recognize diagnostic features. Clearly, visual chunks are used when a quick recognition of the domain is requested. In this experiment, experts demonstrated more ability in recording what was shown in the rock image than novices and they described what they saw using more interpreted features.

Our experimental picture changed completely when another kind of rock was used, such as a metamorphic rock sample in the same short-memory experiments. The result can be seen in Figure 3, where no relation between level of expertise and interpreted features can be identified. The sample here includes a professor of metamorphic petrology, classified (outside his expert field) as a novice. As expected, the geologist exhibited the expert pattern of behavior in his own domain of knowledge. This is evidence that the use of visual chunks is strongly connected with the expertise in the domain.

Another significant result was obtained from the experiments with the set of 9 images shown for one minute each. The experts demonstrated a worse memory for common photographs than the novices, but obtained excellent results when the pictures concerned their domain of expertise. Experts and intermediates could identify the significance of more pictures and more details of each picture, providing a more complete description of each one by comparison with novice performance. In order to describe each picture, the experts have made connections among their constituents using interpretations, such as: "there are two sandstones which experience deep diagenesis; three of marine deposition with dolomitization, and three with shallow diagenesis and good porosity". This demonstrated that the knowledge used for supporting inference was also used to organize and capture information from the domain.

Evidence of ontological support for memorization was also demonstrated when experts were requested to organize other sets of 9 or 7 pictures of common objects (animals and landscapes) and rocks. In the classification of rock pictures, the experts utilized evidently interpreted aspects of the rock, such as the quality of porosity or the kind of cementation. The novices organized the pictures mainly in terms of color, texture or abundance of minerals. When organizing common objects, the experts and novices applied the same geometric characteristics (color, size and shape of animals, and kinds of objects presented in the landscapes).

The organization of the domain in the expert's mind relies on aspects that make the problem-solving process easier or more effective, such as the interpreted features, rather than the taxonomy commonly used for teaching of students. The taxonomy is mainly used to organize the domain for the socialization of knowledge (i.e., for its communication through the community); this is why taxonomies are the first kinds of information collected in interviews for knowledge acquisition. The more the expert understands that a description will be used for someone else, the more he/she will describe it (orally or written) using taxonomic terms. The greater the pressure to provide a good interpretation of the rock, the more interpreted features or real but unnamed (e.g. because tacit) features he/she will use.

We consider here that interpreted features constitute a kind of abstract pattern developed by an expert through experience. It is clear that experts use more of these abstract features than we can collect through experiments like those in our present set, especially because many of the features that support interpretation are not identified by (any) name. We have seen no evidence that interpreted features are interpreted from the image after some steps of inference (short or not). Rather, they are collected through a pattern-matching and instantaneous (more precisely, far too rapid for us to measure) process, which means that they correspond to some internal pattern not directly connected with any object detectable by some low-level or detailed local pattern-recognition process applied to the image. In effect, this is simply one more indication of the importance of tacit knowledge to experts.

5 Rock interpretation modeling in the *PetroGrapher* project

Having dealt with knowledge acquisition and the main differences between geological novices and experts, we now discuss what happens when actual petrographic knowledge is processed in our current system.

According to the results obtained in the cognitive experiments and following the directions of the CommonKADS methodology, the knowledge model of the petrographic application had to take into account the following types of knowledge:

- declarative knowledge applied for communication among people in the domain, also called domain ontology: this includes the taxonomy of the domain, describing mineral constituents, petrographic textures and structures;
- abstract representation of images that support the inference process, the *visual chunks*, as exemplified by Figures 5 and 7;
- interpretation objects, or the objects that are instantiated by the system as a result of the inference process. They are the objects: *Constituent Total, Classification* and *Interpretation*, which latter is, in its turn, an aggregation of *Provenance*, *Diagenetic Environment* and *Diagenetic Sequence.* These objects are shown in the upper-right part of Figure 4 associated with *Sample*, the most general object in the domain.
- the connections of interpretation objects with the visual chunks, here represented by knowledge graphs, as shown in the example of Figure 6;
- the rock interpretation process, described as a PSM, described in Figure 8.

Each type of knowledge, the knowledge acquisition strategy to capture it, and its appropriate representation formalism, are described below.

5.1 The domain ontology

The vocabulary and terms applied by geologists for communication were collected initially through interviews. Interviews provide a better comprehension of the domain, although they were not enough by themselves to define which objects were really necessary for problem solving.

A better result was achieved when a case-based approach was applied. 102 expert descriptions of samples were collected, 12 of them exhaustive and 90 were quantitative descriptions, and their vocabulary and meaning of terms were analyzed. The result was a *partonomy* of terms, in which the terms of the domain are organized primarily through the relation *part-of.* The taxonomy of the domain, though easily found in textbooks, has little utility for realistic petrography, so we decided not to represent it in the knowledge base. (This was a specific choice governed by pattern-matching issues; it should not be understood as general advice for geological knowledge modeling).

Figure 4 shows the objects of the model whose instances are input by the user during the description task (upper left part of picture), the objects that are instantiated by the system by inference (upper right part of picture), and the objects whose instances are part of the knowledge base (right bottom part). This figure can be regarded as a general structure of the ontology.

Visual chunks are pieces of factual knowledge. Although they are not concepts in the domain ontology (in the sense that they are not a shared concept for the geological community), our experience is nevertheless that they are an essential part of the domain model. An expert recognizes sets of features in a specimen or image and then perceives these as a monolithic pattern or a single object. As a result of our experience, we suggest here that we can produce an approximation of this pattern as a visual chunk: an aggregation of rock concepts representing visual geometric features described with the terms of the ontology. This means that a visual chunk connects the expert-level knowledge - the tacit knowledge represented as visual chunks - to novice-level knowledge, represented by the simplified rock concepts associated with geometric features, to make possible the inference steps in the interpretation method. Figure 5 shows how the combination of rock features seen by an expert in an image of a rock can lead to a rock interpretation.

5.2 Evidence-hypothesis relationships

A knowledge graph (K-graph) can be understood as a schema for rock interpretations that describes the many ways in which items of evidence can be used to indicate an interpretation. K-graphs have greater expressivity and bigger granularity when compared to other formalisms that associate items of evidence with hypotheses, such as production rules or Bayesian nets. The K-graph can be described as a tree where:

a) the root node represents the interpretation hypothesis;

b) the leaf nodes represent visual chunks identified by the experts in the image of rock as pieces of evidence necessary to support the interpretation.

Each interpretation is associated in the K-graph with a threshold value that represents the minimum amount of evidence needed to indicate it. The leaf nodes can be combined to increase the influence and the certainty of the interpretation stated.

In our acquisition of knowledge, we focused particularly on acquiring the expert's evidence-hypothesis relations. The expert was provided with the full list of possible diagenetic interpretations, collected from the analyzed cases. Then he was requested to explain what evidence supported each of the listed interpretations. The expert indicated under the microscope many items of evidence and gave them names or short descriptions. For example, the "Continental Meteoric Eodiagenesis Under Dry Climate Conditions" interpretation hypothesis was modeled as a root node, requiring an interpretation measure certainty (a threshold) of 6. This K-graph is presented in Figure 6. It has 7 nodes (each of our graphs has 5 to 7 nodes) to represent the following items: "Iron oxide/ hydroxides", "Infiltrated Clays", "Calcrete", "Dolomite", "Silcrete", "Sulphate", "Dolocrete". The "Calcrete" nodo is detailed in Figure 7.

Next, the list of valid items of evidence that the expert had given was compared with the list of terms in the ontology. Most of the items were not there. It was thus clear that the objects that support inference were not present among the ontological terms used for communication or explanation of the domain.

A new knowledge acquisition approach was then defined to connect the meaning of visual chunks with the objects in the ontology. The expert was provided with his own descriptions on one side and the diagenetic interpretation on other, and requested to connect the terms of the descriptions with parts of interpretation that these terms indicated, using knowledge graphs. The terms were grouped together, as an AND/OR tree of rock concepts describing visual rock features. These rock concepts were then grouped together with AND operators, forming an AND relationship. OR operators were also permitted, meaning that at least one rock feature would need to be found in a matching analysis process. The tree describes how rock concepts are logically combined, to form evidence for an interpretation. This explicit adaptation process can connect basic geometric rock concepts at the user level to abstract concepts at the expert level, to make possible the inference steps in the interpretation method. Our visual chunks modeled in K-graphs were associated with significance indices (weights). For example, the nodes on the "Continental Meteoric Eodiagenesis Under Dry Climate Conditions" K-graph have visual chunks/weights such as: "Silcrete" (weight 6), "Sulphate" (weight 6), "Calcrete" (weight 5) (Figure 7). An interpretation can only become valid when the combination of significance indices of observed visual chunks has reached the threshold associated with it.

We have modeled 6 K-graphs, expressing 6 possible interpretations of the "diagenetic environment in petroleum-reservoir rocks": Continental Meteoric Eodiagenesis Under Dry Climate (detailed in Figure 6); Deep Burial Diagenesis (Deep mesodiagenesis); Eodiagenesis Under Marine Conditions; Early Eodiagenesis Under Mixing Between Marine and Meteoric Waters; Telodiagenesis Under Meteoric Conditions and Continental Meteoric Eodiagenesis Under Wet Climate Conditions. The first of these is shown in Figure 6 and partially detailed below, as an example.

5.3 The problem-solving method for rock interpretation

The dynamic part of knowledge about rock interpretation is described through a PSM designed by (Silva 2001) to model the symbolic pattern-matching process of finding geological evidence to support some interpretation. Moreover, when the PSM (an abstract structure) is particularized to describe specifically a sedimentary geologist's problem-solving activity, it gives us the ability to find the terms in the rock description (stated in a user's vocabulary) that would match with terms in the expert's knowledge and vocabulary for the supporting evidence.

The rock interpretation PSM (e.g. Figure 8) is composed of (1) knowledge roles, which are entities that associate visual chunks with the objects of the domain ontology, (2) individual steps of inference that accomplish a single action over an object of the model, by transforming it, comparing it with other objects, selecting instances or altering its attribute values, and (3) the expected flow of inference.

The rock interpretation PSM has knowledge roles (abstract references to the objects of the model) defined as K-graphs, visual chunks, image description and solutions. The inference steps are: *select* applied to specify visual chunks or K-graphs; *match* used to compare visual chunks with rock features in the problem; and *specify*, which finds possible solutions supported by the visual knowledge. The rock interpretation PSM is able to specify solutions from an image description of rock matched through the use of visual chunks as shown in Figure 8. This basic competence notion is complemented by declarations about what domain knowledge is required by the method. These declarations are called "assumptions" in the technical vocabulary of PSM. Assumptions are used to make explicit, for example, the association of K-graphs and visual chunks. Any association of visual chunks for interpretations relies on the availability of relevance indices (e.g. weights and thresholds) between them, which can model a way of increasing the certainty of solutions and the belief in a rock interpretation.

The rock interpretation PSM is mainly guided by the matching of visual chunks against an image description. But additional search tactics can be used to increase the certainty of an interpretation hypothesis indicated by a K-graph and supported by the visual chunks that are matched. An image is described through basic concept-attribute-value triples (C-A-Vs) of an instance of a rock sample, e.g. in Figure 7, the concept *Diagenetic Composition* has the attribute *Location* with the value *intergranular displacive*. The reasoning can be developed either driven by hypothesis (starting with a selected K-graph) or by data (starting from some observed evidence). In a data-driven inference, a set of visual chunks is logically matched against the C-A-Vs mentioned in an image description. The visual chunks that match are considered as “activated”. Visual chunks are only activated when a minimal logical set of C-A-Vs is found in an image description. The activated visual chunks are used to select other K-graphs that include the same visual chunk, which in their turn are used to select new visual chunks, with the aim of confirming a hypothesis of interpretation. This process is repeated for as long as K-graphs and visual chunks remain to be analyzed, and while visual chunks remain to be matched, over the entire knowledge base. A solution is offered if some K-graph has sufficient confirmation indicated by a set of activated chunks, considering the weights and threshold values. The solutions comprise the matched K-graphs and the whole set of visual chunks activated.

Using the “Continental Meteoric Eodiagenesis Under Dry Climate Conditions” K-graph described in Figure 7 as an example, the visual chunks “Silcrete” and “Sulphate” can be selected initially (the visual chunks with the largest weights). The set of logical combinations of C-A-Vs of these visual chunks is then matched against the image description. If one of these visual chunks is activated, the other chunks “Calcrete”, “Dolocrete”, “Dolomite”, etc, are also selected. This process can increase the certainty about this K-graph currently taken as hypothesis. The new set of selected visual chunks is analyzed, and the certainty of this K-graph is thereby strengthened. Finally, the

"Continental Meteoric Eodiagenesis Under Dry Climate Conditions" solution is indicated by the confirmed K-graph and by the activated visual chunks. As exemplified for this K-graph, the inference process is also repeated over the 6 others "diagenetic environments" modeled, according to their respective K-graphs and visual chunks.

Each of the graphs is tested against the description in an independent way and more than one diagenetic environment can be inferred for the same rock. There is no geological inconsistency in this method, since a rock can experience many different diagenetic environments during the geological time of consolidation. Each process will imprint its results or characteristics over the sediments and will be recognized later through petrographic features. The number of features related to some environment is an indication about how strong was the influence of this environment in the formation of the rock. This is calculated by the inference mechanism, by summing the weights of the chunks that were identified and checking whether the sum is bigger the K-graph threshold. The diagenetic environment (or environments) that the *PetroGrapher* system infers here from its inputs is quite a complex interpretation. If presented in some clear way in conjunction with data from geophysics, it should be very useful in helping geologists to understand the spatial behavior of large-scale petroleum reservoirs in terms of permeability and porosity quality.

Our knowledge acquisition method for eliciting problem-solving behavior has used a *speak aloud* technique, where the experts are observed during the solution task and required to narrate their reasoning during the task. We had to collect and interpret more than 10 such descriptions before the sequence of inference steps became clear. The repeated process was then abstracted, keeping only the main structure of reasoning and the objects of the model that are manipulated by the inferences. In this activity, we were helped greatly by the library of different kinds of PSM provided by the CommonKADS methodology. The model named *Assessment* in that library was used as the starting-point for our derivation of a PSM for *Interpretation*.

6 Some general lessons from the project

In representing and using geological knowledge, the *PetroGrapher* project goes beyond the traditional approach of simple rule-based "expert systems", e.g. because the knowledge is intrinsically too complex for one single simple treatment. We amplify this point in section 6.1 by comparing *PetroGrapher* with one such traditional system.

Part of the complexity arises from the nature of expertise in sedimentary petrography. In section 6.2, we discuss what we have learned about its nature from our study of novice and expert geologists.

6.1 A short comparison with PROSPECTOR

The most familiar achievements in knowledge acquisition for geological domains are the development of the PROSPECTOR system (Duda, Hart et al. 1976; Reboh 1981), which actually influenced the development of the entire subject of knowledge engineering; the XEOD knowledge-based system to identify detrital depositional environments (Schultz, Fang et al. 1988); the Dipmeter Advisor knowledge-based system, for well log interpretation (Smith 1984); and the knowledge acquisition project SISYPHUS III, concerning igneous petrography (Gappa and

Puppe 1998). The common aspect of all these projects was that they have drawn attention to the importance of specific treatment of *visual* diagnostic features and also have shown the necessity of representing the objects that support inference as a separate ontology.

In the PROSPECTOR project, the knowledge was organized in two distinct blocks: knowledge bases for special and general purposes (Reboh 1981). The former represents the geological models of each kind of known ore deposit described as inference networks. An inference network expresses the relation among items of geological evidence (in PROSPECTOR, most items are descriptions of visual features) and hypotheses of interpretation. Several uncertainty indices are associated with the network, indicating how strongly each piece of evidence influences the conclusion, and the prior probability of occurrence of any relevant ore deposit is taken into account using the Bayesian method (Duda, Hart et al. 1976).

The general knowledge base represents the ontology of the domain, described through the taxonomy of the geological terms. Assertions of geological facts about the domain are also part of this knowledge base. Reboh (1981, page 11) describes them as corresponding "to nodes in network 'in a higher level', as inference networks of specific purpose". In their analysis, the authors admit the necessity of two cognitive levels of knowledge representation, although they were not able, at that time, to provide an adequate solution to deal with these objects. The higher level of objects was described through *partitioned semantic networks.* These networks aggregate some terms of the taxonomy (mixing concepts and instances of concepts from the knowledge base) and describe their associations, in order to represent the semantic content of objects that, when are found together, provide special support for inference – which we describe as visual chunks. We believe that the authors of PROSPECTOR, as in the other projects mentioned here, have identified the difficulties in acquiring and representing visual chunks, but without providing an adequate treatment in terms of knowledge acquisition and knowledge representation such as we have developed in our project.

The *PetroGrapher* system provides a more elegant solution: a knowledge base that is easy to update and maintain, and which recognizes distinct cognitive objects and represents those using different formalisms. Also, in making explicit the problem-solving method that, in the PROSPECTOR system, is hidden in the programming code, it can implement a genuinely knowledge-independent system, whose knowledge can be updated and reused during the natural evolution of the domain.

6.2 Classification, diagnosis or assessment: what do petrographers do?

The work in cognitive psychology that underlies CommonKADS pays special attention to the detailed differences among the various PSMs that experts in different fields use. (Gardner, Rush et al. 1998) contains examples of both the variety of PSMs and the details of what goes into a PSM: it lists and describes about 30 PSMs, which certainly do not exhaust the range of possibilities. For successful knowledge engineering, it is necessary to avoid mistakes in the identification of the best PSM(s) to use as the framework for eliciting and expressing any particular kind of expert knowledge. This is not easy for geological expertise, because there are at least 3 types of candidate: classification, diagnosis, and assessment.

According to (Schreiber, Akkermans et al. 2000), classification is concerned with establishing the correct class for an object using the object characteristics, e.g. (in geology) classification of minerals in a rock. Following (Clancey 1985), the heuristic classification problem is concerned with abstract observations, which are sometimes used in place of simplified observations to generate hypotheses. Here, too, rock interpretation is mainly based on qualitative abstractions (sets of visual factors associated with each other and understood together) taken as observations, which can be used to specify interpretations. But rock interpretation is not a class of some well-defined scheme of interpretation; nevertheless, we can still find solutions supported by visual aspects that can produce a weaker outcome, e.g. a merely "acceptable" interpretation.

According to (Benjamins and Jansweijer 1994), the goal of diagnosis is to find solutions that explain both the initial and any additional observations. Diagnostic PSMs are firstly based on fault models, or models of abnormal behavior (mainly defined from fault diagnosis in technical systems). Data about the abnormal behavior guides the reasoning and search to reach the solution of the problem. In contrast to that, the rock interpretation process tries to explain some observations, but there are no fault models to support it. Furthermore, the observations are not monotonic data in the interpretation process, but more abstract information, here defined as visual chunks. Rock interpretation also raises issues about management of certainty, which has been covered extensively in research on medical expert systems but which is nevertheless only weakly described in diagnostic methods.

The assessment problem described in (Schreiber, Akkermans et al. 2000) prompts the proposal of a category for a case, based on a set of domain-specific norms. The rock interpretation PSM can be characterized as a something very similar, because it is based on a case of image description and norms modeled as K-graphs and visual chunks. The challenge for a rock interpretation PSM with respect to the assessment method is to consider explicitly an *image-based* reasoning process of interpretation, establishing the assumptions and requirements that arise from visual knowledge modeling. Understanding the chunking process as a simple analysis of interpretation norms also does not take into account many cognitive aspects that are involved in this process of interpretation. We expect to be able to explore these aspects by trials of our newly-derived PSMs on further instances of petroleum-reservoir interpretation.

In our view, assessment is a better PSM match (even though not a perfect one) than classification or diagnosis to the details of what we observe when experts in sedimentary petrography solve problems. We would be interested in further observational evidence from other experiments similar to our own, especially if they advance alternative views of what petrographers actually do.

7. Remarks on the implementation of *PetroGrapher*

The requisites of the *PetroGrapher* system were to define a complete standard in petrographic descriptions of siliciclastic rocks, providing a helpful support in the use of terminology, scales and compositional classification methods. In addition, the system is intended to support the knowledge management functionalities: helping to capture the expert knowledge in petroleum reservoir evaluation, providing information access for those who need the results of petrographic analysis, and making it easy to combine the information with other knowledge sources. These requirements have led us to the definition of a three-layer intelligent database application for petrographic descriptions.

The system provides a full environment for siliciclastic rock description, where the user can describe a rock qualitatively with a predefined vocabulary and format, and make a quantitative description with the support of a special hardware appliance - the electronic stage - connected to the microscope. This interface component allows the capture of the petrographic information that is stored in a relational database system, through the mapping of knowledge objects to entities in the database model. Independent inference components can be applied to the data, providing the following information: compositional name of the rock, in the terms given by 6 different authors and methods; probable provenance of the sediments in the terms given by McBride (McBride 1963); diagenetic environments that have determined the rock consolidation (whose PSM is described in the present paper) and diagenetic sequence, based on the paragenetic relation of constituents, according to (De Ros 1994). All the results are stored in the database along with the description and can be exported in XML format for further use. The whole set of descriptions can be consulted by a navigational tool based on the ontological terms or through an SQL standard interface.

The interpretation PSM is the most complex component of inference. In knowledge-based systems, reasoning patterns can be implemented using different software components, such as logic systems, numerical systems, symbolic systems, database systems, Web components, etc. Below, we detail the implementation of the rock interpretation PSM using a software architecture for interpretation, which is a specification of knowledge and data software components. The challenge for this architecture is the integration of symbolic and database components, dealing with storage and querying over a large volume of knowledge and data, in addition to the reasoning resources for interpretation. The components of the architecture are:

a) interpretation engine – which translates inference steps into knowledge/data requests, and carries out the task of parsing symbolic knowledge or data structures used to represent the inference steps of the rock interpretation PSM. This engine also has the role of analyzing the inference knowledge, such as visual chunk weights and interpretation thresholds, according to the interpretation criteria;

b) query engine – which supplies the interaction between the interpretation engine and the database components, writing and running queries automatically, according to the knowledge/data requests that are sent by the interpretation engine;

c) database components – supported by a relational database system, which has schemata for a repository of rock descriptions and a repository of inferential knowledge about rock interpretation, here modeled as K-graphs and visual chunks.

This architecture is also supported by i) an image-description model mapped to an entity-relationship model, producing a database schema to store rock descriptions and ii) inferential knowledge (K-graph and visual chunk models) mapped to an entity-relationship model, producing a database schema for storage of the concepts and relationships of the inferential knowledge used in the rock interpretation.

The operation of the "interpretation architecture" (Figure 8) is as follows. When an interpretation request is received (I), the interpretation engine loads the mapping between knowledge and data models (II). The working memories (WMs) are organized as symbolic data structures according to these mappings, because the inferential knowledge is stored in the main memory during the interpretation. The query engine also employs these mappings, using them to produce queries (in SQL), which allow the integration of the knowledge and data components defined in the architecture. Next, the query engine receives knowledge requests (III). These requests are parsed and knowledge queries (IV) are serviced in the database of inferential knowledge. This process is repeated until the query engine has completely loaded every knowledge structure of interpretation in the knowledge WM (V - VIII). Then, the interpretation engine parses these structures (VI), starting a sequence of data requests to the query engine (VII). The query engine parses the data requests and runs the resulting description queries on the database of application data (IX). The application data about rock samples are gradually queried and stored in the data WM by the query engine (X - XI). The interpretation engine parses these data about rock samples (XII) regarding the interpretation structures stored in the knowledge WM (VI), trying to match an interpretation indicated by K-graphs and visual chunks. The overall interpretation process is controlled by a symbolic algorithm of interpretation, which is implemented in the interpretation engine. This algorithm can run the inference steps described in the rock interpretation PSM, managing the main requirements of the process of image-based interpretation. Finally, if an interpretation has been reached by the analysis of inferential knowledge modeled, it will be returned by the interpretation component (XIII).

A more complete description of *PetroGrapher*'s resources and system architecture is given in (Abel, Silva et al. 2004).

8 Concluding remarks

We have inspected and recorded the cognitive analysis of the expert when solving geological problems. Differently from many other scientific domains, geology requires from professionals not only the intensive use of knowledge, but also abilities in building abstract three-dimensional dynamic models and telling stories about something that happened millions of years ago. These abilities were observed, in this project, through the investigation of expertise in sedimentary petrography.

Geologists work by recognizing visible geological features in nature and providing an interpretation about their overall geometry and their processes of formation. During the evolution from novices to experts, geologists develop their ability by training their low-level cognitive visual system and by building symbolic abstractions of images to interpret the acquired data. We call these structures visual chunks.

The main cognitive structures that can be identified in the solution of problems of geological interpretation are:

- the set of visual chunks and associated geological meanings, which is acquired;

- the set of schemata, which builds the ontology of the domain, whose role is abstracting the common features and organizing geological concepts, allowing a common comprehension and communication among the community of geologists.

During the learning process, a novice turns into an expert basically by acquiring a bigger number of visual chunks and associating them with the concepts of the domain ontology, and by developing more efficient and sophisticated methods for their application. Geologists apply these knowledge structures through an investigation process initially guided by the visual chunks matched against the task and later by the knowledge embodied within the ontology and its structure They reach their final conclusions by compounding the partial meanings to arrive at a complete interpretation, subject to the general rules about the domain.

On the most general level, the study of expertise provides us with a better comprehension about the abilities required to deal with geological tasks and allows us to determine the properties required of an effective knowledge acquisition process for the development of knowledge-based systems. We have found that the combined use of knowledge graphs and case analysis is the best way of making explicit and acquiring the declarative knowledge and causal relations of the domain, which were not evident in elicitation sessions conducted on a traditional knowledge-acquisition basis. These cognitive resources integrate the collection of implicit expert knowledge in geological reservoir characterization.

Our conclusions about knowledge engineering are testable through the performance of the *PetroGrapher* system, which has now been introduced into a real corporate environment. Its behavior, as monitored there by an expert, is encouraging: no wrong conclusions, and less than 15% of interpretations missing from complete descriptions of specimens. Further, it is about to be used in staff training, which will test the realism and cognitive validity of the ontology that has been derived.

Acknowledgements

This work is supported by the Brazilian Council of Research – CNPq, with the special program CTPETRO, Project No. 502009/2003-9. We thank Laura Mastella for her many constructive suggestions and thorough reviewing of the first version of this paper.

References

Abel, M., 2001. The study of expertise in Sedimentary Petrography and its significance for knowledge engineering. Ph.D. Thesis (in Portuguese: "Estudo da perícia em petrografia sedimentar e sua importância para a engenharia de conhecimento"), Programa de PG em Ciência da Computação, Porto Alegre, Universidade Federal do Rio Grande do Sul.

Abel, M., Castilho, J. M. V. and Campbell, J., 1998. Analysis of expertise for implementing geological expert systems. World Conference in Expert Systems, Mexico City, Cognizant Communication Offices, pp. 170-177.

Abel, M., Silva, L. A. L., De Ros, L. F., Mastella, L. S., Campbell, J. A. and Novello, T., 2004. *PetroGrapher*: managing petrographic data and knowledge using an intelligent database application. Expert Systems with Applications 26(1 SPECISS), 9-18.

Angele, J. and Lausen, G., 2004. Ontologies in f-logic. Handbook on ontologies. S. Staab and R. Studer. Berlin, Springer-Verlag:, pp. 29-50.

Antoniou, G. and Harmelen, F. v., 2004. Web ontology language: OWL. Handbook on Ontologies. S. Staab and R. Studer. Berlin, Springer-Verlag, pp. 67-92.

Baader, F., Calvanese, D., McGuinness, D., Nardi, D. and Patel-Schneider, P., 2003. The description logic handbook. Theory, implementation and applications. Cambridge, Cambridge University Press.

Benjamins, R. and Jansweijer, W., 1994. Toward a competence theory of diagnosis. IEEE Expert 9(5Oct): 43-52.

Borst, W. N., 1997. Construction of engineering ontologies. Enschede, University of Twente.

Brickley, D. and Guha, R. V., 2000. Resource description framework (rdf) schema specification 1.0., W3C Proposed Recomendation.

Buchanan, B. and Shortliffe, E., 1984. Rule-based expert systems: the MYCIN experiments. Reading, Addison-Wesley.

Clancey, W. J. 1985. Heuristic classification. Artificial Intelligence. Amsterdam. 27, pp.289-350.

De Groot, A. D. 1965. Thought and choice in chess. Mouton, The Hague.

De Ros, L. F., Morad, S. and Paim, P. S. G. (1994). The role of detrital composition and climate on the diagenetic evolution of continental molasses: evidence from the Cambro-Ordovician Guaritas Sequence, southern Brazil. Sedim. Geol. 92: 197-228.

Duda, R. O., Hart, P. E., Barret, P., Gaschnig, J., Konolige, K., Reboh, R. and Slocum, J., 1978. Development of the PROSPECTOR consultation system for mineral exploration. Menlo Park, Stanford Research Institute International.

Duda, R. O., Hart, P. E. and Nilsson, N. J., 1976. Subjective Bayesian methods for rule-based inference systems. SRI Project 4763. Palo Alto, CA, Stanford Research Institute International, pp. 1075-1082.

Ericsson, K. A. and Smith, J., 1991. Toward a general theory of expertise: prospects and limits. New York, Cambridge University Press.

Fikes, R. and Kehler, T., 1985. The role of frame-based representation in reasoning. Communications of the ACM 238(9), 904-920.

Gappa, U. and Puppe, F., 1998. A study of knowledge acquisition - experiences from the SISYPHUS III experiment for rock classification. Workshop on Knowledge Acquisition, Modeling and Management, Alberta, Canada, Voyager Inn.

Gardner, K. M., Rush, A., Crist, M., Konitzer, R. and Teegarden, B., 1998. Cognitive patterns: problem-solving framework for object technology. Cambridge, Cambridge University Press.

Gómez-Pérez, A. and Benjamins, V. R., 1999. Overview of knowledge sharing and reuse components: Ontologies and problem-solving methods. International Joint Conference on Artificial Intelligence(IJCAI-99), Workshop on Ontologies and Problem-Solving Methods (KRR5). V. R. Benjamins, B. Chandrasekaran, A. Gomez-Perez, N. Guarino and M. Uschold. Stockolm, Sweden.

Harmon, P. and Sawyer, B., 1990. Creating expert-systems for business and industry. New York, John Wiley &Sons.

Kelly, G. A., 1955. The psychology of personal constructs. New York:, Norton.

Leake, D. B., Ed., 1996. Case-based reasoning: experiences, lessons and future directions. Cambridge, MA, AAAI Press/MIT Press.

Leão, B. F. and Rocha, A. F., 1990. Proposed methodology for knowledge acquisition: a study on congenital heart disease diagnosis. Methods of Information in Medicine(29): 30-40.

McBride, E. F., 1963. A classification of common sandstones. Journal of Sedimentary Petrology 33, 664-669.

Newell, A., 1982. The knowledge level. Artificial Intelligence. Amsterdam. 18, 87-127.

Nonaka, I., Takeuchi, H. and Takeuchi, H., 1995. The knowledge-creating company: how Japanese companies create the dynamics of innovation. New York, Oxford University Press.

Polanyi, M., 1966. The tacit dimension. Knowledge in organizations. L. Prusak. Newton, MA, Butterworth-Heinemann, pp. 135-146.

Puerta, A. R., Egar, J. W., Tu, S. W. and Musen, M. A., 1992. A multiple-method knowledge-acquisition shell for the automatic generation of knowledge-acquisition tools. Knowledge Acquisition. [S.l.]. 4, 171-196.

Reboh, R., 1981. Knowledge engineering techniques and tools in the PROSPECTOR environment. Palo Alto, CA, Stanford Research Institute International.

Smith, R. G. (1984). On the development of commercial expert systems. AI Magazine. [S.l]. **5:** 61-73.

Schreiber, G., Akkermans, H., Anjewierden, A., Hoog, R. d., Shadbolt, N., Velde, W. v. d. and Wielinga, B., 2000. Knowledge engineering and management. Cambridge, The MIT Press.

Schultz, A. W., Fang, J. H., Burston, M. R., Chen, H. C. and Reynolds, S., 1988. XEOD: an expert system for determining clastic depositional environments. Geobyte. [S.l]. 3, 22-32.

Silva, L. A. L., 2001 The use of problem-solving methods in rock interpretation tasks, Master Thesis (In Portuguese: Aplicando métodos de solução de problemas em tarefas de interpretação de rochas). Programa de PG em Ciência da Computação. Porto Alegre, UFRGS.

Sowa, J. F., 1984. Conceptual structures: information processing in mind and machine. Reading, Addison Wesley.

Studer, R., Benjamins, V. R. and Fensel, D., 1998. Knowledge engineering: principles and methods. Data & Knowledge Engineering. Amsterdam. 25, 161-197.

Velde, W. v. d., 1993. Issues in knowledge level modeling. Second generation expert systems. J. M. David, J. P. Krivine and R. Simmons. Berlin, Springer Verlag.

Wright, G. and Ayton, P., 1987. Eliciting and modeling expert knowledge. North Holland, Elsevier Science Publ.

Corresponding author:
Mara Abel, Instituto de Informática - UFRGS, Av. Bento Gonçalves, 9500 Bloco IV. 91.501-970 Porto Alegre, Brazil. Tel.: +55 (51) 33166825, Fax: +55 (51) 33167308, E-mail address: {marabel}@inf.ufrgs.br.