Negotiating a Shared Conceptual Model using Groupware

Debbie Richards
Department of Computing
Division of Information and Communication Sciences
richards@ics.mq.edu.au

Abstract

This paper describes an approach that begins with asynchronously capturing individual conceptual models into separate knowledge-based systems using a technique known as Ripple-Down Rules. The rules are combined and used by Formal Concept Analysis to develop the group's conceptual model thereby revealing common ground and differences between the stakeholders. As group process is offered which allows conflicts to be identified and resolved, where possible, using our negotiation strategies and resolution operators. The individual models are updated based on the negotiations and the cycle repeats. The end result over a number of iterations is a shared conceptual model.

Keywords

AL01 Knowledge representation; AL04 Knowledge acquisition; HB26 Simulation and modelling IS; Formal Concept Analysis; Group Decision Support; Ripple Down Rules

1. INTRODUCTION

The motivation for this work is based on research into acquiring knowledge into knowledge-based systems (KBS[1]). It has become well accepted that the bottleneck in the development of KBS is acquiring the knowledge from experts who typically have difficulty articulating the knowledge they use. Similar to the findings for other groupware systems, problems magnify and additional issues emerge when we move from a single user to considering knowledge from multiple sources. Gaines and Shaw (1989) found that experts disagreed with other experts within their own domain and even with themselves over time. The approach proposed here addresses the issues of knowledge acquisition (KA) from individual experts and the identification and reconciliation of conflicts between these individuals. The approach and tool we developed is designed for individual and group use. Unlike some groupware whose only goal is to assist the group communication process, our approach also results in a number of artifacts which are used to generate new concepts and refine those already identified. Initially, the tool is used by individual experts to develop their own knowledge base (KB), which they own and control. The user is not constrained in the terminology they use although the terms used in the rule conditions tend to be features of the case as the KA technique we used is case-based. The next section gives a brief introduction to our KA technique. The individual KBs are combined (the rules are literally thrown in together) and used to automatically generate a concept lattice. The combined conceptual model is used to provide feedthrough to individuals and the group.

[1] The term KBS is used synonomously with expert systems (ES) in this paper.

Explanation has been a key feature of KBS. The ability to provide the reasoning process behind a decision, commonly in the form of a rule trace, was seen as a major benefit not offered by other types of systems. We use the explanations offered when traversing rule pathways to assist in the negotiation process between multiple stakeholders. With the use of a technique known as Formal Concept Analysis (FCA) (Wille 1982, 1992) the rule pathways are restructured and organized into a hierarchy using term subsumption.

It would be difficult to classify our approach and tool using any one of the types of systems given in (Dix et. al 1998). Our tool can be seen as a meeting and decision support system in so far as it allows individuals to record their reasoning (arguments) when used to build their own KB and to support the discussion of ideas and concepts when used in face-to-face groups that are synchronously co-located. What makes our approach different to typical meeting and decision support systems is that the team members work at times alone and at other times together to develop individual as well as a co-authored system. The shared KB provides structure, focus and makes identification of similarities and differences within the group providing a wider communication bandwidth not available when reviewing the individual KBs separately.

2. THE UNDERLYING THEORIES

The knowledge acquisition and representation technique known as Ripple Down Rules (RDR) (Compton and Jansen 1990) is a hybrid case-based and rule-based approach which supports rapid and incremental acquisition of knowledge by the domain expert. We wanted an approach which allows the expert to acquire and maintain their own KBS in keeping with the needs of users, even truer of domain experts, to own and control their own knowledge (Langlotz and Shortliffe 1983, Kidd and Sharpe1987). From an individual's point of view, RDR supports easy KA, which is performed in a reflexive mode. This means that minimal analysis or reflective thought is required. The expert simply looks at a case, assigns a new conclusion if they disagree with the one assigned by the system, and then identifies one or more attribute-value pairs which justify new conclusion. We also hold a socially-situated and evolving view of knowledge (Clancey 1997) which is supported by the RDR paradigm. Through the use of cases and the Multiple Classification RDR (MCRDR) (Kang, Compton and Preston 1995) exception structure, knowledge is patched locally in the context in which it was acquired. In this way the cases provide grounding (Clark and Schaefer 1989) in the real world. The use of cases and their associations to the rules is invaluable when it comes to group decision support as they let the rule or concept owner identify a situation in which their knowledge is considered valid which may then be argued with the group. The offering of counterexamples is a technique that can assist KA (Wille 1989) but coming up with a counterexample during discussion is often difficult for humans.

The MCRDR rules are used to generate a concept lattice using FCA. The conditions in the rules consist of attribute-value pairs from the case and thus can be considered to be primitives. Higher concepts are found and organized into a complete lattice by finding the intersections of shared rule conditions. A concept in FCA is seen as a set of attributes and the set of objects that share those attributes. The lattice makes explicit the implicit structure, relationships and abstractions which are so difficult for experts to describe. The concept lattice allows the user to reflect on and explore the knowledge in the KB. The individual may generate lattices for their own KB for validation and explanation purposes. The individual KBs are used to develop a lattice which shows the combined knowledge of the experts. Unlike a co-authored system, these individual KBS are retained and updated as determined by the individual. The shared model is regenerated each time from the updated individual KB. Of course, some of the changes will be due to discussion between the experts over the shared conceptual model. As will be described in more detail later, in the shared model it is possible to tag certain concepts to be excluded and to reconcile differences in terminology via a subsumes table.

To reduce the complexity of the lattice and information overload and to improve clarity and response times, lattices may be derived based on various selection criteria. These criteria are described in (Richards 2000). A detailed description of the RDR KA technique and representation or the mathematics underlying FCA are not given in this paper. Interested readers are directed to (Richards and Compton 1997) for this information. We look further at the group process in the next section.

3. A GROUP SCENARIO

As mentioned, our approach begins with individuals developing their own KB which are then combined to produce a shared conceptual model. This section describes the activities that occur during the group phases. A possible scenario is presented to describe these activities and how they may lead to a model which is representative of the group. The scenario is based on some data from the SISYPHUS III (Shadbolt 1996) experiment which included knowledge on the classification of 19 well-known igneous rocks from multiple sources of expertise. Unless the reader is familiar with the domain of geology, this example may be difficult to follow. However, this is typical of any domain of expertise and a reason why we use RDR and FCA that allows domain experts to directly enter and reconcile their own knowledge. It is hoped that the reader will gain a feel for the process rather than understand why certain decisions were made. The process assumes that the group of domain experts have some familiarity with reading a concept lattice and that there is a group facilitator managing group interactions. The SISYPHUS III experiment, our results and a comprehensive discussion of the reconciliation process from a more technical point of view is given in (Richards and Menzies 1998).

3.1 Concept Comparison and Conflict Detection

Once the individual KBS have been identified and a merged concept lattice developed, the group meet to compare the viewpoints and identify conflicts. A number of researchers offer different sets of conflict types (e.g. Easterbrook 1991 and Schwanke and Kaiser 1988). We have adopted the four quadrant model of comparison between experts developed by Gaines and Shaw (1989). This model classifies two conceptual models as being in one of four states:

Consensus is the situation where experts describe the same concepts using the same terminology.
Correspondence occurs where experts describe the same concepts but use different terminology.
Conflict is where different concepts are being described but the same terms are used.
Contrast is where there is no similarity between concepts or the terminology used.

Discovering the consensus between conceptual models establishes common grounds from which differences can be viewed. We generally take a broader view of conflict to encompass inconsistencies that include the states of contrast, correspondence and conflict. Gaines and Shaw's model, however, offers greater precision in describing the nature of the conflict, which is important in deciding how it can be handled. Before we look at conflicts, we need to understand how to read the concept lattice. A concept lattice is shown in Figure 1. The labeling has been reduced for clarity. Each small circle on the diagram represents a concept. Attributes (in our case rule conditions) that belong to a concept are found at the node and by traversing ascending pathways. Objects (identified by rule number – expertise source - conclusion code) are found at the node and by traversing descending pathways. For example, Concept number 6 includes the set of attributes {DARK_MINERALS=LT30, FELDSPAR=1_3TO2-3PLAGIOCLASE, GRAIN_ SIZE=COARSE, 1=1} and the set of objects {2-C3-%AD000}.

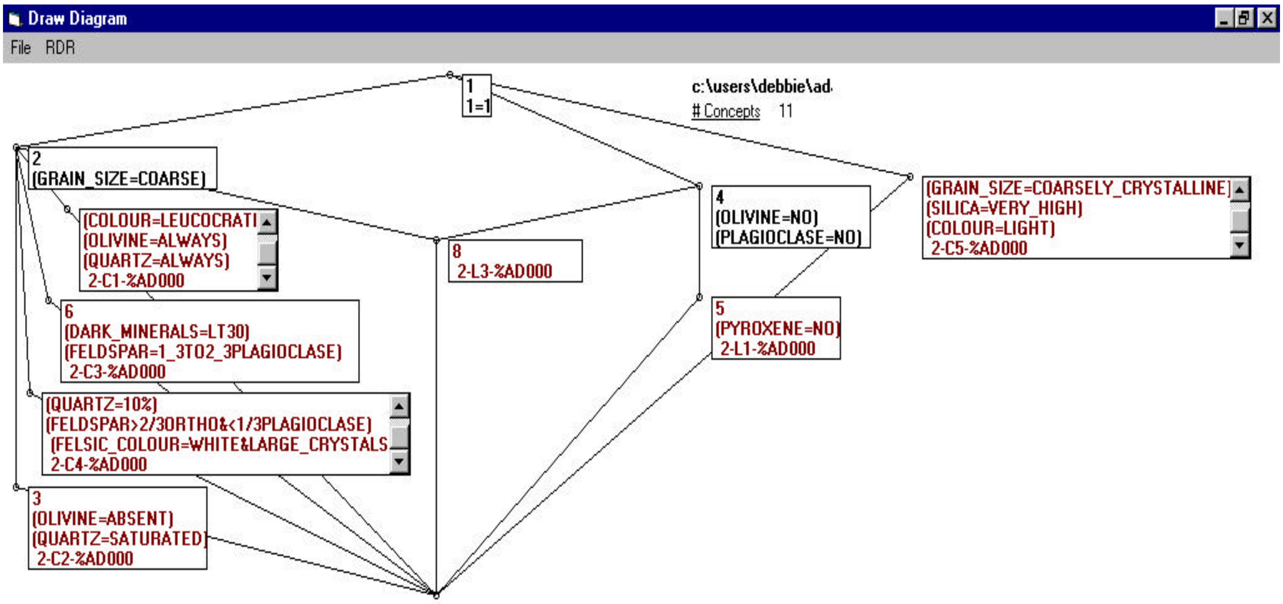

Figure 1: The Concept Lattice for the Conclusion %AD000- Adamellite based on seven KBS.

The concept lattice in figure 1 concerns rules which classify a rock as Adamellite and includes rules from 7 different sources of expertise, identified as C1, C2, C3, C4, C5, L1 and L2. We can analyse the lattice to determine the differences between the expert viewpoints. We can see that the C5 viewpoint (far-right) for Adamellite is in a state of total contrast since none of its attributes are shared with any of the other viewpoints. However, some of these differences appear to be terminology related. There is consensus between C1, C2, C3 and L2 that GRAIN_SIZE=COARSE but in C5 the GRAIN_SIZE= COARSELY_ CRYSTALLINE. This appears to be a correspondence type of conflict. There appear to be other correspondence errors. The attribute QUARTZ is used in C1, C2 and C4 with the values ALWAYS, SATURATED and 10%, respectively. The value of OLIVINE in L1 and L2 is NO and in C2 the value is ABSENT. In C5 the COLOUR=LIGHT and in C1 the COLOUR= LEUCRATIC. The dictionary meaning of “leuco” is white (Macquarie Dictionary), so it appears that the terms in these two concepts are compatible. It also appears that DARK_MINERAL =LT30 also indicates a lightness of colour. The differences in the terminology used for the values assigned to GRAIN_SIZE, QUARTZ, OLIVINE and COLOUR can be reconciled by using the synonym table to map to a common term as shown in Table 1.

The value assigned to OLIVINE in C1 is ALWAYS and represents a conflict where the terms are compatible but the concept is obviously the opposite to the concepts in L1, L2 and C2. There is consensus between L1 and L2 that PLAGIOCLASE = NO but these concepts conflict with the concepts FELDSPAR=1_3TO2_3 PLAGIOCLASE for C3 and FELDSPAR>2/3ORTHO &<1/3PLAGIOCLASE for C4. These various conflicts need to be resolved which takes us to our next stage.

3.2 Conflict Negotiation

Before we can decide how to fix a detected inconsistency we need to provide a conflict resolution strategy. There are a number of resolution methods including negotiation, arbitration, coercion and education (Strauss 1978). Negotiation is the most appropriate within the assumed context of parties of equal status and ability. A good solution will require creativity and creativity is not something that can be automated. However, since automation is a fundamental goal of this project we extend our approach beyond a general, genial chat by offering as much automated assistance for this step as possible.

A number of resolution strategies have been offered (e.g. Easterbrook 1991, Thomas 1976). Easterbrook and Nuseibeh (1996) offer five categories that covers the actions we have found necessary. These are:

1. Resolve, correct any errors;
2. Ignore, no action is performed;
3. Delay, identify the existence of the inconsistency but defer action until a later date;
4. Circumvent, identify the existence of the inconsistency so it can be avoided;
5. Ameliorate, reduce the degree of inconsistency. This action requires analysis and reasoning.

Resolving conflict will involve performing modifications. If the cause of disagreement is differences in terminology (correspondence in the Gaines and Shaw four state model) then one technique is to up-date all views to conform to an agreed upon set of terminology. This option is probably not satisfactory to the various stakeholders and also means that the history of changes is being lost or altered. A simple and more appropriate solution is to use a synonym table which maps terms from individual views into a shared terminology which are then used for comparison.

Original Term	Synonym
GRAIN_SIZE=COARSLY_ CRYSTALLINE	GRAIN_SIZE=COARSE
QUARTZ=SATURATED	QUARTZ=ALWAYS
QUARTZ=10%	QUARTZ=ALWAYS
OLIVINE=ABSENT	OLIVINE=NO
COLOUR=LEUCRATIC	COLOUR=LIGHT
FELSIC_COLOUR= WHITE&LARGE CRYTALS	COLOUR=LIGHT

Table 1: The Synonym Table

Another way in which conflict may be resolved is through the addition or deletion of attributes or objects. By changing rule conditions or the conclusion, concepts can be brought into a state of consensus.

The last four resolution strategies are relevant for situations in which a complete resolution cannot be negotiated and each one has its appropriate usage. For example, ignoring is a useful strategy where the issue is not that important or pursuing it is not worth the effort or harm it may cause to the end solution. These approaches can be termed as living with inconsistency or 'lazy' consistency (Narayanasway and Goldman 1992) and can be compared to fault-tolerant systems that continue to function after non-critical failures occur. We also accept that living with inconsistency will be necessary and use tags to identify the status of the conflict. The use of tags is similar to the use of "pollution markers" (Balzer 1991) that act as a warning that code may be unstable or that the users should carefully check the output. Pollution markers can be used to screen inconsistent data from critical paths that must have completely consistent input. If it is the concept that is being circumvented, ignored or delayed, we mark the concept in the shared model since there is not necessarily a one-to-one correspondence between rules and concepts.

The updated shared model and updated individual KBs are used as input in the generation of the next shared model.

In addition to terminological differences we have some conceptual differences. To demonstrate how the approach can be used to resolve the difference, the following scenario simulates some of the discussion and decisions that could occur amongst the group members.

- Since sources C2, L1 and L2 agree that OLIVINE=NO, the expert in C1 realises that he has made an error and changes OLIVINE=ALWAYS to OLIVINE=NO. The individual KBS is updated to reflect this change. This also requires amending the value of OLIVINE in the associated hypothetical case or passing a case from another viewpoint which covers this situation..

- The L1 expert agrees that GRAINSIZE=COARSE should be included so he amends his.

- A feature of our approach is the ability to offer counterexamples that can be used in negotiations. In Figure 2 the case associated with Concept No. 8 (Rule 2 in Card Sort 5) is shown to the group to assist with reconciliation of this conflict. The C5 and L1 experts can not be persuaded by the other experts to drop the attributes SILICA=VERY_HIGH and PYROXENE=NO, respectively, so it is decided to delay resolution of these conflict until another meeting. This is achieved by using the Delay Tag. When a concept is delayed, the control background and foreground colour is reversed and the word DELAYED is displayed. Figure 2 shows concepts 8 and 9 have been delayed. The user may also identify that a concept is not worth further consideration by using the Ignore Tag which drops the concept from view.

- A final strategy concerns the handling of the controversy over the importance of FELDSPAR in determining if a rock is Adamellite. Expert C3 believes that the FELDSPAR content is one to two-thirds PLAGIO-CLASE. Expert C4 believes FELDSPAR is only less than one-third PLAGIO-CLASE, experts L1 and L2 believe that PLAGIOCLASE = NO and experts C1 and C2 do not consider FELDSPAR or the PLAGIOCLASE content. It is thus decided to circumvent the concepts with these attributes. This is achieved by tagging Concepts No 10, 11 and 13 in Figure 1 as circumvented. Once given this tag a concept is not included in determination of the list of predecessors (parents) and list of successors (children) which are used to layout the line diagram. It can be seen in Fig 2 that these attributes are no longer shown. If desired, these concepts can be reinstated and shown.

All of the changes mentioned above are reflected in our final diagram in Figure 2. Note that even though the number of concepts has only reduced by 1 the concepts are much less complex. In figure 1 GRAIN_SIZE= COARSE offered the most, but not total, point of agreement. Now all views agree with this and there are more attributes shared by viewpoints that previously only appeared in one viewpoint. As shown visually in 2, the viewpoints in Card Sorts 1,2,3, and 5 are more similar to each other than the viewpoints in Laddered Grids 1 and 3 which are similar to each other.

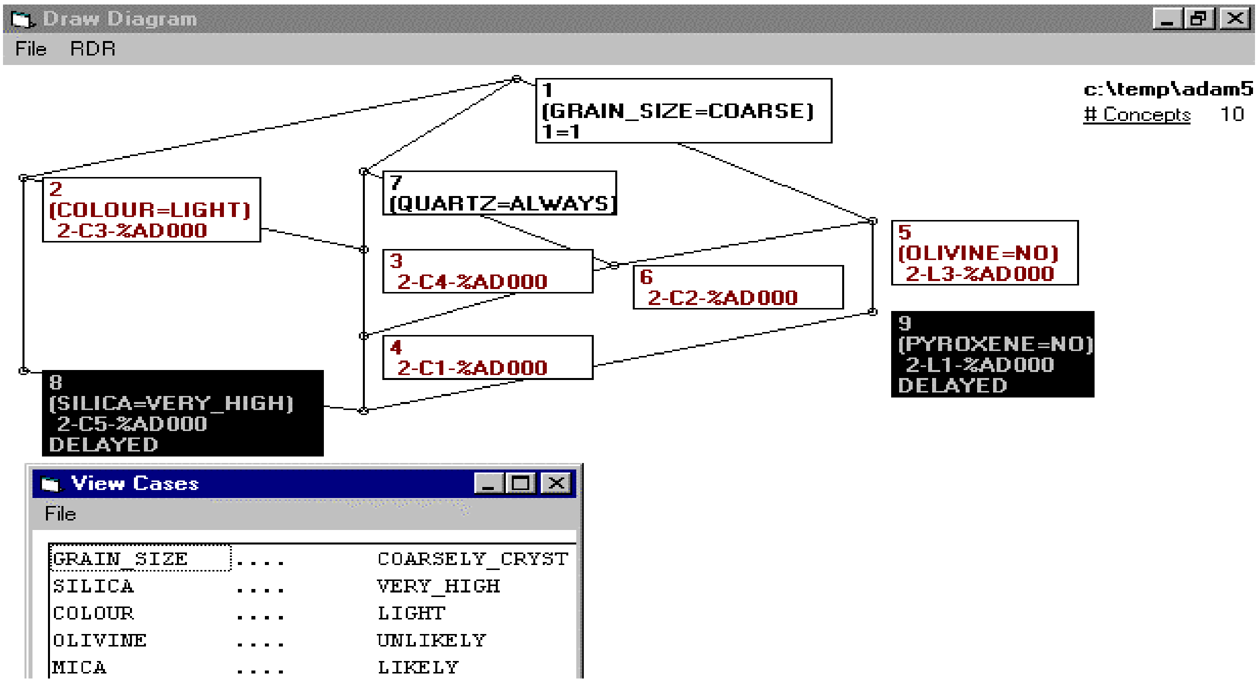

Figure 2: The final Line Diagram screen from this round of negotiations. Some attributes have been dropped or added to views, concepts have been tagged to be circumvented (not shown) or delayed (shown). The case for concept 8 is shown to assist with negotiation. There is considerably less conflict now than in Figure 1.

3.3 Repeating the Individual-Group Cycle

To determine that our RE strategy is resolving conflict we need to employ some measures of the degree of conflict before and after. By computing a score for each concept in each viewpoint compared to each other viewpoint and taking the total of these scores we can check that the degree of conflict after the RE process is less than at the start. We assign a score of 0 to a concept found to be in a state of consensus with a concept in another viewpoint, since the distance between them is zero. For concepts in a state of conflict we take the number of attributes (conditions) that they have but do not share divided by the total number of attributes. This assumes that the two concepts share the same object (conclusion). If they do not then it appears that they are not meant to represent the same concept so that comparison is not meaningful. For concepts in a state of contrast (no partial or complete match in the other viewpoint) we assign a score of 1, which is the same result as if we used the conflict measure since the number of attributes not shared divided by the number of attributes is equal to one. Concepts in a state of correspondence are treated the same as concepts in conflict since we are ignoring the reason for the differences and are just interested in the size of the difference. Once terminology differences are reconciled such concepts will move into one of the other states and be handled accordingly. Using these measures we computed the degree of conflict at three points, at the beginning, after updating the synonym table and at the end.

Round 1	Start	Synonym	End
C1	4.89	3.25	1.98
C2	4.76	3.02	2.50
C3	4.76	3.52	2.60
C4	4.89	3.25	2.39
C5	6.00	3.35	3.01
L1	5.33	5.00	3.17
L3	4.10	3.48	2.84
TOTAL	34.74	24.88	18.49

Table 2: The degree of conflict between each viewpoint for the %AD000-Adamellite conclusion before RE.

In Table 2 the total amount of conflict at each of these three phases is shown. At the start of the meeting, we see that all viewpoints are in conflict with others, with views C1, C2, C3 and C4 having similar degrees of conflict. From this table, and supported by the lattice in Figure 1, we can see that viewpoint C5 is in complete contrast with all other views with the highest degree of conflict, followed by L1 that only shares some attributes with L3. L3 has the lowest degree of conflict. We also see that the total degree of conflict for all viewpoints has reduced from 34.74 before we began our resolution strategies to 24.88 after we applied our first strategy of reconciling terms. It is interesting to note that all views except L1 now have similar, though lower than before, degrees of conflict. This shows that much of the conflict originally in C5 was due to differences in terminology, which we have already discovered in our previous discussions. Very little of the conflict in L1 appears to be terminology related. After we have applied the remaining resolution strategies we have not removed all conflict but the overall degree of conflict has reduced by 53% from 34.74 to 18.49.

4. EVALUATING THE CONCEPT LATTICE

We have chosen to use a visual representation of the individual and shared models as a central part of our group decision support software. The utitlity of the approach thus hinges on the usefulness and usability of the concept lattice. We have conducted a small survey to evaluate a number of aspects of the line diagram including: how easy it was to learn to read the line diagram; the value of the line diagram representation over a linear textual rule trace; and how well the diagram could be used for learning about a domain and the knowledge in a KBS. The survey found that 10 out of the 12 subjects were able to learn to read a line diagram within a few minutes, that the line diagram was easier and faster to use than a rule trace in answering questions about the knowledge base and that even novices could reason about the knowledge using the line diagram allowing the tool to be used for such purposes as hypothesis testing and tutoring. The results were promising but, as noted by Kremer (1998) and evidenced in (Petre and Green 1993), use of a visual language requires time and effort to learn and this makes evaluation of the line diagram by novices a difficult task.

A different type of evaluation has also been undertaken. Four case studies have been also conducted in four different domains: agriculture, chemistry, geology and pathology. In these studies a domain expert was asked to comment on the knowledge gleaned by a domain beginner (a level lower than novice) when the beginner explored KB's built about these domains by other people. It was found that the concept lattice opened up a channel of communication and opportunity for learning which was not possible by simply looking at the rules or rule traces. The results indicated that the novice had gained an understanding of the key concepts in each domain and that the line diagrams were useful for motivating discussion between the novice and the expert and for identifying missing and erroneous knowledge on the part of the student. The results of the survey with novices and the evaluation with domain experts is not conclusive or large enough to allow statistical sampling. However, as a preliminary evaluation the results are promising. A detailed description and discussion of the survey with novices and case studies with experts is given in (Richards 1998). The generality of the approach to other KR and the value of the concept lattice as a taxonomic and ontological representation have also been evaluated (Richards 1999).

5. DISCUSSION AND FUTURE WORK

The KA and conceptual modeling technique have been well established and used by numerous others both for research and commercially. The usefulness of RDR and FCA have been proven through acceptance by users but virtually no experimental studies have been performed to explore the reasons for their success or potential areas for improvement to users. The absence of such studies reflects the general lack of interest in user-cooperation issues. This shortcoming has been identified for some time (Langlotz and Shortliffe 1983, Salle and Hunter 1990) but little is being done to change the focus of KBS research. The cost of performing experiments with users is high. When these users are experts in their domain, the costs are even higher. However, we believe the lack of interest in HCI issues is not primarily due to cost issues. The difficulties associated with KA has resulted in KBS approaches which have become focused on the reuse and sharing of problem-solving methods and ontologies. These approaches require the system and user to interact via the mediation of a knowledge engineer. The end user has become even further removed from research concerns and the needs of the knowledge engineer have gained prominence. This is very unfortunate as it has been recognized that even where knowledge is found to be accurate and reliable the KBS becomes an undertutilised resource when it is not presented in a way that fits the mental model of the user (Salle and Hunter 1990).

Another difficult issue we are addressing is the reconciliation of conflicting sources of expertise. Many KA research groups are currently focused on the reuse and sharing of ontologies as a means to alleviate the KA bottleneck. Many KA research groups are currently focused on the reuse and sharing of ontologies as a means to alleviate the KA bottleneck. Such endeavours appear to take for granted that a shared view already exists and ignore the problem of reconciling differences. Skuce (1995) sees the main problem is that the underlying ontological and terminological assumptions have not been made sufficiently explicit or agreed upon. Other solutions to the SIS III example given in this paper, ranged from choosing one source to focus on (avoidance) (Erdmann 1998) to replacing the given KA material with an alternative single source, that is, a reference book (abandonment) (Jansen, Schreiber and Weilinga 1998). We did not want to take either extreme and decided to tackle conflict between sources as the focus of our solution. A less extreme approach is to develop some negotiation strategies which expects differences to be sorted out prior to entering the knowledge into the system (*a priori* alignment). This could have involved detecting the conflicts and choosing which knowledge source to accept for each concept in conflict. We rejected the avoidance, abandonment or *a priori* alignment options because we believed that there are a number of good reasons for capturing, tracking and reconciling different viewpoints. By deciding to tackle head on the issue of multiple sources we needed to consider the associated group process to support the reconciliation process. A preliminary group process has been offered in this paper.

Evaluation of the MCRDR/FCA approach is high on our priority list. Our investigations will be in the area of requirements engineering which is particularly affected by soft issues such as human-to-human communication, group dynamics and organizational behavioural factors. We still have some technical (and financial) hurdles to jump but our fundamental interests are supporting and adding rigour to the group process of specifying requirements. We will initially use our tool on a number of case studies. These case studies will use a variety of sources of requirements such as direct elicitation of rules, use cases and data-flow diagrams. The case studies will include well-documented case studies for comparison with other approaches and solutions. The case studies will provide insight into how requirements can be acquired and presented in machine-readable format. These requirements will be used by FCA to develop concept lattices. Data, in the form of concept lattices, observations, and sets of requirements for individual and shared models will be collected at each iteration of the capture – compare – negotiate - reconcile cycle. Experiments will be performed that allow us to evaluate the impact of variables, such as the source of requirements, the format of requirements, the use of different views (these are subsets of the total set of requirements) and so on, on the efficacy of our approach. It is anticipated that our findings from the case studies and experiments will result in modifications to our approach and tool.

Evaluation of the approach will extend beyond feasibility, which can be measured using the case studies and experiments. Further evaluation will focus on the usefulness and usability of the approach. We want to answer questions such as:

- ✓ Can we get people to use the approach?
- ✓ Can the approach not only identify commonalities and conflicts between stakeholders but also assist in resolution of these conflicts so that a representative set of requirements can be developed?

A number of issues related to usablity are still to be resolved. At this stage we envisage a central computer around which a group of stakeholders are seated. As the requirements gathering group is typically made up of people with different roles and responsibilities, including users of the system to be developed, it is envisaged that someone skilled in using the system would sit at the terminal to make appropriate changes. This person or another that was skilled as a group facilitator would be able to guide the meeting by identifying the key points for discussion, suggesting possible resolution strategies and assist the group in gaining a shared understanding as well as manage who had the floor. Since the goal is to produce an artefact that represents the combined viewpoints it is important that such a facilitator be present so that the discussion result in changes that head towards consensus. The ability to use a distance-based metric as described in section 3.3 which is also visually represented in the lattice to show the progress being made is seen as a potentially useful feature.

Evaluation will explore the usefulness of the approach for group decision making and will thus involve comparison with similar groupware tools. We do not propose to begin with a study of such tools since our focus is on the application of an existing KE technique to requirements engineering (RE) and not the application of existing groupware approaches to RE. RE approaches often assume that the requirements are already available in table format. We find this assumption unrealistic and plan to explore the automated conversion of natural language into a crosstable. To this end we intend to explore tools such as ATTEMPTO (Fuchs and Schwitter 1996) which takes in constrained natural language and outputs propositions. We will also explore work done by (Al-Ani et. al. 1999) which uses gIBIS (Conklin and Begeman 1989) to develop requirements from natural language.

Evaluation will be performed with real subjects. We plan to use undergraduate student, postgraduate students in a Postgraduate Professional Development Program and participants in the Software Requirements Engineering (SRE) Mailing List[2] which is comprised of academics, students, researchers and RE practitioners. Various tests (survey, case study or experiment) will need to be designed depending on the method of access; time, experience and knowledge of the subjects; analysis techniques and the goals of the particular test. To allow distributed users to participate in group decision making we would need to address the issue of distributed meeting rooms. We may be able to integrate our approach with the work by Greenberg and Roseman (1998) which uses a room metaphor to allow work to occur individually and as a group as well as synchronously and asynchronously. We need to consider the numerous issues that differentiate face-to-face communication from text-based communication. Some recent findings indicate that the group may more successfully achieve their goals by working in a distributed mode without the emotional complications of face-to-face interaction (Damian et. al. 1999). It is clear that extensive evaluation is a critical next

---

[2] SRE mailing list is an electronic discussion bulletin board that is managed and moderated by Didar Zowghi at Macquarie University. It has close to 700 subscribers world wide.

step in refining and determining the usefulness and usability of this approach and we look forward to being able to report our findings in the not too distant future.

References

Al-Ani, B., Leaney, J., and Lowe, D. (1999) A Taxonomy of Partially Excluded Service Descriptions *Proceedings of the Fourth Australian Conference on Requirements Engineering (ACRE'99)* 29-30 September, 1999, Macquarie University, Sydney, 15-26.

Balzer, R. (1991) Tolerating Inconsistency *Proceedings of 13th International Conference on Software Engineering (ICSE-13)* Austin, Texas, USA, 13-17th May 1991, IEEE Computer Society Press, 158-165

Clancey, W. (1997) The Conceptual Nature of Knowledge, Situations and Activity In Feltovich, Paul J., Ford, Kenneth, M. and Hoffman, Robert R. (eds) *Expertise in Context: Human and Machine* AAAI Press/MIT Press, Cambridge, Massachusetts, 248-291

Clark, H. and Schaefer, E. (1989) Contributing to Discourse *Cognitive Science* 13:259-294

Compton, P. and Jansen, R., (1990) A Philosophical Basis for Knowledge Acquisition. *Knowledge Acquisition* 2:241-257

Compton, P., Edwards, G., Srinivasan, A., Malor, P., Preston, P. Kang B. and L. Lazarus (1992) Ripple-down-rules: Turning Knowledge Acquisition into Knowledge Maintenance *Artificial Intelligence in Medicine (*4) 47-59

Conklin, J. and Begeman, M. gIBIS: A Tool for all Reasons *Journal of the American Society for Information Science,* March 1991.

Damian, D., Eberlein, A. Shaw, M. and Gaines, B. (1999) Studies in Distributed Software Requirements Engineering In *Proceedings of the 12th Workshop on Knowledge Acquisition, Modeling and Management* (KAW'99), 16-21 October, 1999, Banff.

Dix, A., Finlay, J., Gregory, A. and Russell, B. (1998) *Human Computer Interaction* 2nd Edition, Prentice Hall, London.

Easterbrook, Steve (1991) Elicitation of Requirements from Multiple Perspectives *PhD Thesis*, Department of Computing, Imperial College of Science, Technology and Medicine, University of London, London SW7 2BZ.

Easterbrook, S. and Nuseibeh, B. (1996) Using Viewpoints for Inconsistency Management *BCSEEE Software Engineering Journal* January 1996:31-43.

Erdmann, Michael (1998) Formal Concept Analysis to Learn from the Sisyphus-III Material *11th Workshop on Knowledge Acquisition, Modeling and Management*, Banff, Canada, SRDG Publications, Departments of Computer Science, University of Calgary, Calgary, Canada, Vol 1:SIS-2.

Fuchs, N.E. and Schwitter, R. (1996) Attempto Controlled English (ACE) In Proceedings of the First International Workshop on Controlled Language Applications (CLAW'96) Leuren, Belgium.

Gaines, B. R. and Shaw, M.L.G. (1989) Comparing the Conceptual Systems of Experts *The 11th International Joint Conference on Artificial Intelligence* :633-638.

Greenberg, S. and Roseman, M. (1998) Using a Room Metaphor to Ease Transitions in Groupware *Research Report 98/611/02* Department of Computer Science, University of Calgary, Calgary, Alberta, Canada, January, http://cpsc.ucalgary.ca/grouplab/papers/

Jansen, M.G., Schreiber A. Th. and Weilinga, B.J. Rocky III - Round 1 A Progress Report In Gaines, B. and Musen, M. (eds) *Eleventh Workshop on Knowledge Acquisition Modeling and Management (KAW'98)* Banff, Canada, April 18-23, 1998, SIS-5.

Kang, B., Compton, P. and Preston, P (1995) Multiple Classification Ripple Down Rules: Evaluation and Possibilities *Proceedings 9th Banff Knowledge Acquisition for Knowledge Based Systems Workshop* Banff. Feb 26 - March 3 1995, Vol 1: 17.1-17.20.

Kidd, A. L. and Sharpe, W. P. (1987) Goals for Expert System Research: An Analysis of Tasks and Domains In *Expert System IV* Cambridge Uni. Press, 146-152.

Kremer, Rob (1998) Visual Languages for Knowledge Representation Eleventh Workshop on Knowledge Acquisition Modeling and Management (KAW'98) Banff, Canada, April 18-23, 1998.

Langlotz, Curtis P. and Shortliffe, Edward E. (1983) Adapting a Consultation System to Critique User Plans. *Int.l Journal of Man-Machine Studies 19*:479-496.

Narayanaswamy, K. and Goldman, N. (1992) “Lazy Consistency”: A Basis for Cooperative Software Development *Proceedings of International Conference on Computer-Supported Cooperative Work (CSCW’92)* Toronto, Ontario, Canada, 31 October- 4 November, 257-264; ACM SIGCHI & SIGOIS.

Petre, M. and Green, T.R.G. (1993) Learning to Read Graphics: Some Evidence that ‘Seeing’ an Information Display is an Acquired Skill Journal of Visual Languages and Computing 4(1):55-70.

Richards, D. (1998) The Reuse of Knowledge in Ripple Down Rules Knowledge Based Systems *PhD Thesis,* University of New South Wales.

Richards, D. (1999) Visualising Knowledge-Based Systems *Proceedings of the Third Workshop on Software Visualisation*, Dec 3-4, 1999, University of Technology, Sydney, 1-8.

Richards, D. (2000) The Visualisation of Multiple Views to Support Knowledge Reuse *Intelligent Information Processing* (IIP’2000) An international conference in conjunction with the 16th IFIP World Computer Congress WCC2000, August 21-25, 2000, Beijing, China (to appear)

Richards, D. and Compton, P. (1997) Uncovering the Conceptual Models in Ripple Down Rules In Dickson Lukose, Harry Delugach, Marry Keeler, Leroy Searle, and John F. Sowa, (Eds) (1997), *Proceedings of the Fifth International Conference on Conceptual Structures (ICCS'97),* August 3 - 8, University of Washington, Seattle, USA, Lecture Notes in Artificial Intelligence, Springer-Verlag, Number 1257, Berlin pp:198-212

Richards, D. and Menzies, T. (1998) Extending the SISYPHUS III Experiment from a Knowledge Engineering to a Requirements Engineering Task In Gaines, B. and Musen, M. (eds) *11th Workshop on Knowledge Acquisition, Modeling and Management (KAW’98),* Banff, Canada, SRDG Publications, Departments of Computer Science, University of Calgary, Calgary, Canada, Vol 1:SIS-6.

Salle, J.M. and Hunter, J., (1990) Computer/User Cooperation Issues for Knowledge-Based Systems: A Review, *Technical Report AUCS/TR9003*, Aberdeen University

Schwanke, R.W. and Kaiser, G.E. (1988) Living with Inconsistency in Large Systems *Proceedings of the International Workshop on Software Version and Configuration Control* Grassau, Germany, 27-29 January 1988, 98-118;B.G. Teubner, Stuttgart.

Shadbolt, N., (1996)URL:http://www.psyc.nott. ac.uk/aigr/research/ka/SisIII

Skuce, D. (1995) Conventions for Reaching Agreement on Shared Ontologies *Proceedings of the 9th Banff Knowledge Acquisition Based Systems Workshop,* Banff, 3-1 – 3-19.

Strauss, A. (1978) *Negotiation: Varieties, Contexts, Processes and Social Order* Jossey-Bass Publishers, San Francisco, CA.

Thomas, K. (1976) Conflict and Conflict Management In Duneette (ed) *Handbook of Industrial and Organisational Psychology* Rand McNally College Publishing Co.

Wille, R. (1982) Restructuring Lattice Theory: An Approach Based on Hierarchies of Concepts In I. Rival (ed*), Ordered Sets* Reidel, Dordrecht, Boston, 445-470.

Wille, R. (1989) Knowledge Acquisition by Methods of Formal Concept Analysis In E. Diday (ed), Data Analysis, Learning Symbolic and Numeric Knowledge Nova Science Pub., New York, 365-380.

Wille, R. (1992) Concept Lattices and Conceptual Knowledge Systems *Computers Math. Applic. (23*) 6-9:493-515