

Acquiring and Applying Contextualised Tacit Knowledge

*Debbie Richards and Peter Busch
Department of Computing,
Division of Information and Communication Sciences,
Macquarie University, Sydney
Ph: +61 02 9850 8567
Fax: +61 -2 9850 9551
{richards, busch}@ics.mq.edu.au*

Abstract

The acquisition and application of knowledge, in particular tacit knowledge, are seen as decisive competitive factors in the knowledge society of the twenty-first century. Despite much talk about the importance of knowledge transfer, little research shows how to identify and measure tacit knowledge, less research addresses how to transfer tacit knowledge between individuals and even fewer of these approaches offer any technology that can assist with transference. The approach outlined in this paper, known as Ripple Down Rules (RDR), is not concerned with identifying who has tacit knowledge but with how to capture tacit knowledge from those identified by some means as experts. Unlike most knowledge acquisition, the RDR knowledge acquisition technique does not rely on the expert to specify what they know. Instead, tacit knowledge becomes codified by the RDR system while the domain expert exercises his/her expertise. The knowledge may be transferred to another individual through our recent extensions to RDR which uses Formal Concept Analysis to retrospectively and automatically develop knowledge models that the user can explore. The high degree of participation, ownership and control afforded by the RDR technique together with the simplicity of the approach enables and encourages user satisfaction and utilisation of the system. Others within the organisation can apply that knowledge by executing the rules or they can automatically generate and compare models to internalise that knowledge. In either case, the result is that the knowledge stays in the organisation after the individual expert retires or leaves. This addresses two knowledge management challenges: utilisation and preservation of knowledge.

1. The Acquisition and Application of Knowledge and the Role of Context

The acquisition and application of knowledge are seen as “the decisive competitive factors” in the knowledge society of the twenty-first century (van Daal, de Haas and Weggeman, 1998, p. 255). However, decades of research by the expert systems/knowledge based systems (KBS) community has taught us a number of lessons. Firstly, the acquisition of knowledge turned out to be a very difficult task forcing review of what knowledge or expertise really is and questioning the very notion of knowledge transfer. Knowledge transfer was based on the physical symbol system hypothesis (Newell and Simon, 1976) and treated knowledge as some ‘stuff’ that existed in an expert’s head and which could be transferred to a machine. In the early 90s it became apparent that the notion of knowledge transfer was naïve and was replaced by a focus on modelling knowledge at a level above its symbolic representation, known as the knowledge level (Newell, 1982).

It is interesting to observe that the term knowledge transfer is currently popular in the knowledge management (KM) literature. It is important that the term be understood to mean transfer between individuals rather than the earlier notion of transfer from human to machine. If we take into account the socially situated nature of knowledge (Clancey 1997), this new understanding of knowledge transfer is more realistic. However, despite all the talk about the necessity of knowledge transfer, little research exists that even begins to identify and measure tacit knowledge (TK) within an organisation. The works of Sternberg (1999), Reber (1993), Reed and Hock (1983), Busch and Richards (2000) are among the few. Even less research offers a way of transferring knowledge between individuals.

The difficulties associated with capturing expertise led to knowledge being defined as either explicit or tacit (Nonaka, Takeuchi and Umemoto 1996), although overwhelming the majority of research to

date has focussed on the explicit component of knowledge. Current KBS research is predominantly concerned with the development of ontologies as a way of acquiring domain and task structures. The previous focus was on the development of general problem solving methods (PSM). Even though ontological and PSM research are founded on knowledge-level modelling, the problem with both of these foci is that they again imply that knowledge is an artefact that humans are capable of articulating. Also, the type of knowledge being acquired is primarily codifiable knowledge. Tacit knowledge is often treated as that knowledge which can't be captured. However, research has shown that there is a component of tacit knowledge that can become articulated (Busch and Richards 2000; Grant and Gregory 1997; Raghuram 1996; Nonaka, Takeuchi and Umemoto 1996; Howells 1995; Goldman 1990; Pylyshyn 1981). Further justification for pursuing tacit knowledge acquisition comes from research in the growing field of knowledge management where it has been found that an organization cannot afford to leave tacit knowledge in the too-hard basket.

Another lesson to be learnt from KBS research is the importance of context on the applicability of knowledge. The sharing and reuse of knowledge has been the driving force of the 90s in KBS research as a means of addressing the difficulties associated with the initial capture of knowledge. The desire to share and reuse knowledge led to identification of the need to capture the context as well as the knowledge (e.g. Chandrasekaran and Johnson 1993, Guha and Lenat 1990, McCarthy 1991 and Patil *et al* 1992) so that the knowledge can be adapted to fit the new situation (Clancey 1992). The socially situated view of knowledge places even greater emphasis on the role of context. Situated cognition involves taking into account interaction between the individual's inner state and the external environment and trying to record all the influencing factors. In addition, since thinking and acting interact by modifying each other, context is socially situated and affected by such things as activities, participation, roles, contribution and norms (Clancey, personal communication). The situated view rejects the notion, which some hold (e.g. Noh *et al.* 2000), that knowledge, including tacit knowledge, is stored in memory and simply needs to be retrieved in the appropriate circumstances. Instead, knowledge is seen to evolve and to be "made-up" to fit each situation. Thus, a situated view of knowledge places great emphasis on incremental techniques that allow change, capture context and which acquire knowledge without relying on a human to state or codify that knowledge.

The approach outlined in this paper, known as Ripple Down Rules (RDR) (Compton and Jansen 1990), is based on a situated view of knowledge and offers a way of capturing tacit knowledge. The RDR knowledge acquisition technique does not rely on the expert to specify what they know. Instead, knowledge is captured while the domain expert exercises his/her expertise. The domain expert is not asked to develop models of the domain or to offer explanations of their reasoning processes. RDR performs codification of the tacit knowledge. The knowledge may be transferred to another individual through our recent extensions to RDR which allow retrospective and automatic development of knowledge models that the user can explore. For model generation we make use of Formal Concept Analysis (FCA) (Wille 1992). FCA takes the RDR as input and generates a set of concepts which are ordered into a complete lattice. When lattices from multiple experts are combined (Richards 2000a) the resulting lattice can be viewed as an ontology because the lattice provides a specification of a shared conceptualisation (Gruber 1993). Figure 1 shows the process of tacit knowledge management using the RDR/FCA approach. These steps are described in later sections.

The RDR technique is a hybrid case-based and rule-based approach. Context is an important aspect of RDR, which is captured in associated cases and the exception structure. Context is also critical in the FCA technique and captured in what is known as a formal context. In the next section we introduce RDR and discuss the relationship between cases and rules and our handling of context. In Section 3 we introduce context in FCA and how it is being used with RDR. In Section 4, we relate our work in KBS to current work within the KM field. Closing remarks appear in Section 5.

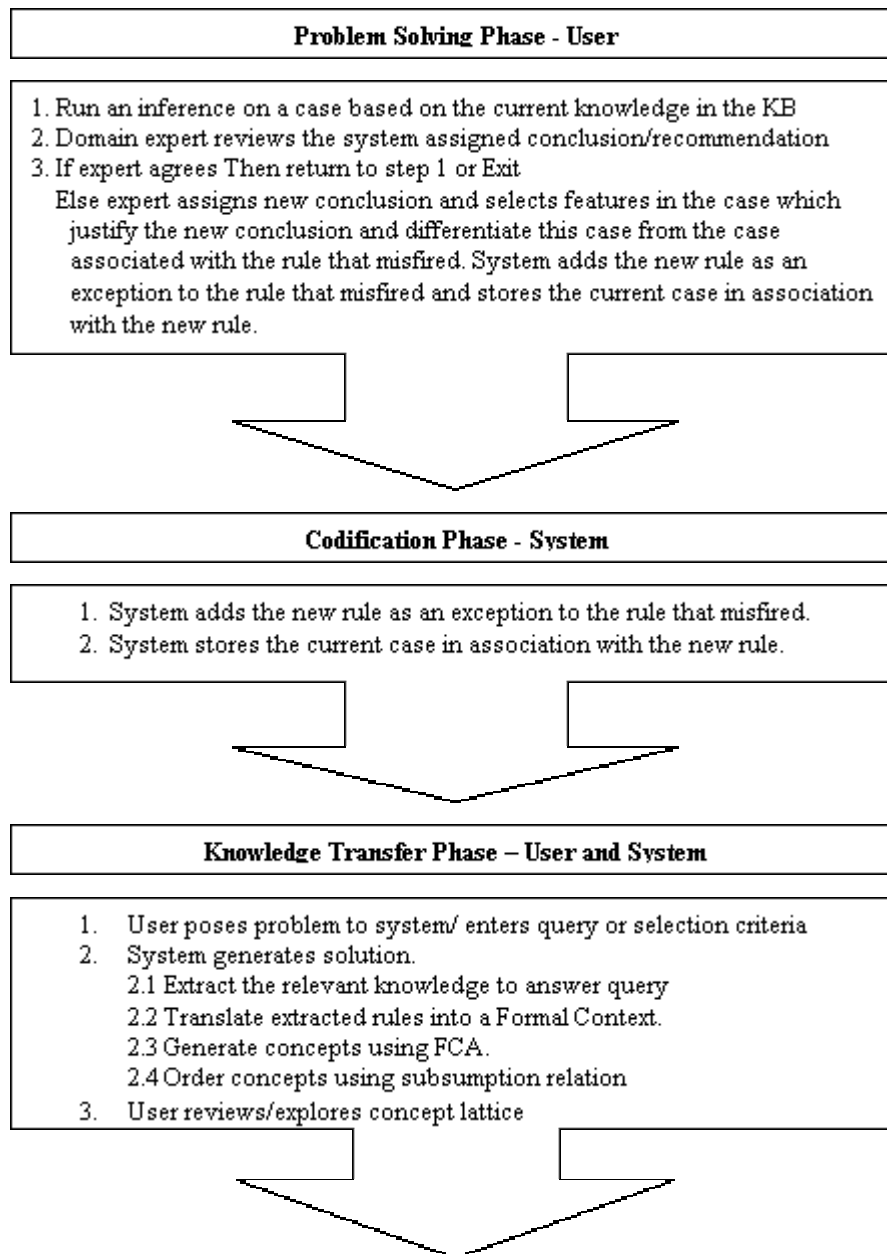


Figure1. The Knowledge Management Process

2. What are Ripple Down Rules ?

RDR were created to address the problem of maintenance of large KBS and in the belief that knowledge is not an artefact which only needs to be properly defined in order to be used. Instead the recommendation given by an expert depends on the context in which it is given and does not consist of a description of the expert's thought processes but is a justification of why that recommendation was made (Compton and Jansen 1990). The context consists of the case and the person receiving the knowledge. A number of variations of RDR have been developed. The first implementation was known as single classification RDR and went into routine use in a large Sydney hospital performing pathology report interpretation (Edwards *et al.* 1993). This paper will only consider the more recent implementation known as multiple classification RDR (MCRDR).

2.1 Multiple Classification RDR

MCRDR were developed to handle classification tasks where multiple independent classifications are required (Kang, Compton and Preston 1995, Kang 1996). This method builds n-ary trees and consists only of exception branches. A better description may be sets of decision lists joined by exceptions. Figure 2 shows an example MCRDR showing two levels of decision lists. An MCRDR is defined as the quadruple $\langle \text{rule}, P, C, S \rangle$, where P is the parent rule, C are the children/exception rules and S are the sibling rules within the same level of decision list. Every rule in the first list is evaluated. If a rule evaluates to false then no further lists attached to that rule are examined. If a rule evaluates to true all rules in the next list are tested. The list of every true rule is processed in this way. The last true rule on each path constitutes the conclusions given. Current empirical evaluations of a commercial system using MCRDR have shown that experts can build systems with 3- 4,000 rules in about one person week.

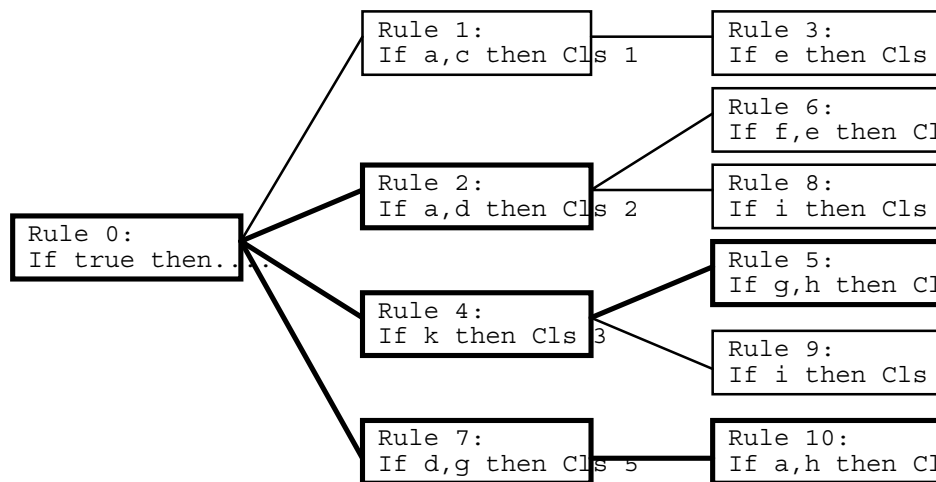


Figure 2. An MCRDR KBS.

The highlighted boxes represent rules that are satisfied for the case $\{a,d,g,h,k\}$. We can see that there are three conclusions, Class 2 (Rule 2), Class 5 (Rule 5) and Class 8 (Rule 10).

As shown in the problem-solving phase in Figure 1, knowledge acquisition in MCRDR involves the expert performing an inference on a case. In response the system will provide one or more conclusions. The expert reviews each recommendation. If the expert disagrees with any conclusion they select that conclusion and enter or select the correct conclusion. The expert then selects one or more features (attribute-value pairs) in the case which justify the new conclusion. To assist in selection of features and to provide online validation, the expert is shown a case associated with the rule that misfired, known as a cornerstone case, which must be distinguished from the current case. Multiple cases may be associated with the one rule. The expert reviews each case selecting features to form an exception rule. If any cases in the cornerstone list are not covered by the new rule, the user is shown that case and must add a feature from that case to the rule until all related cornerstone cases are distinguished. Remarkably, the expert provides a sufficiently precise rule after two or three cases have been seen (Kang, Compton and Preston 1995). Once the user has finished specifying the rule, the system performs the codification phase shown in Figure 1. The user does not need to know where to add the rule. The user may decide to stop an incorrect conclusion instead of replacing it with a new conclusion. This is achieved by adding a stopping rule which has a null conclusion in the same way. Simulation studies (Kang, Compton and Preston 1995) have shown MCRDR to be a superior representation to the original RDR structure by producing knowledge bases that mature more quickly and are more compact even for single classification domains. It is conjectured that this occurs because more use is made of expertise rather than depending on the knowledge base structure (Kang 1996). The process of reviewing a case and assigning a conclusion to it and picking features that justify why one case should have one conclusion and another case should have a different conclusion is what experts are good at doing and part of their routine work.

2.2 Using Cases to Provide the Context of Rule-Based Knowledge

We can view a case as “a contextualised piece of knowledge representing an experience that teaches a lesson fundamental to achieving the goal of the reasoner” (Kolodner 1993, p.13). From this definition, the two key aspects of a case are its ability to teach one or more lessons and specification of the context in which the lesson(s) apply. In case-based reasoning (CBR), the capturing of context is done by the assignment of indices to cases. As with RDR, the “purpose of indexing is to differentiate a case” (Kolodner 1993, p.266). The founders of RDR saw the benefits of using cases and also the benefits of rule-based systems. However, both approaches had weaknesses. Rule based systems did not provide context and grounding in the real world. CBR suffers from the indexing problem associated with retrieving appropriate cases and deciding how best to adapt cases. RDR attempts to draw on the strengths and to avoid the limitations of both.

The use of cases in RDR is similar to the use of cases in CBR. Cases are used to:

- assist the user to develop rules,
- provide the appropriate context of the new rule being entered through storage of cornerstone case and
- validate the entered rule by ensuring the current case is differentiated from other cases associated with the incorrect rule.

When the expert picks the salient features in the case (these form the rule conditions) they are in effect selecting the index by which to retrieve the case. The creation of indexes is one of the major problems in CBR because most approaches try to generate these automatically. Getting the expert to provide this index as a natural part of their duties substantially simplifies the task. The way that a ripple-down rule is used to differentiate between cases combines the approaches of difference and checklist-based indexing and the failure-driven nature of learning in RDR makes it similar to explanation-based indexing. See Kolodner (1993) for description of these three main approaches to indexing cases. While RDR does place importance on the role of cases, each case provides a local context and there is no attempt to define generally applicable rules or landmark cases. CBR techniques usually try to identify important cases that are used as significant cases which can be used to classify new cases.

3. Applying Formal Concept Analysis to Ripple Down Rules

We have recently applied ideas from FCA to MCRDR so that the implicit concepts, both primitive and more abstract, in an MCRDR assertional KBS can be found and structured into a terminological KBS. FCA supports the knowledge transfer phase shown in Figure 1. The RDR assertions elicited and codified and in phases 1 and 2, respectively, provide a performance system and the concept lattice derived using FCA in phase 3 provides an explanation system. This enhancement is important as minimal analysis of domain knowledge was a strength of the KA technique (Compton *et al* 1993) but it meant that it was not possible to show higher level models of the domain knowledge. The abstraction hierarchy that FCA automatically develops offers a retrospective model that is based on the MCRDR rules. While the MCRDR approach does not require a model for KA or inferencing we are interested in providing the user with a model since models have been found to be beneficial for instruction (Schon 1987) and for explanation (Clancey 1993) which are important for knowledge transfer.

RDR and FCA agree that the domain expert should be the one responsible for directly performing KA. They also share a number of views regarding knowledge. This is particularly true of the role of context. FCA is:

“guided by the conviction that human thinking and communication always take place in contexts which determine the specific meaning of the concepts used” (Wille 1996, p. 23).

RDR and FCA do not consider the knowledge captured to be globally applicable but relevant within the given context. This notion has been formalized into what is known as a Formal Context and is shown as a crosstable as in Figure 6. The crosstable is used to find concepts. A concept in FCA is

comprised of a set of attributes and a set of objects that share those attributes. A formal context is a triple (G,M,I) where G (for Gegenstande in German) is the set of objects which forms the extension of the concept, M (for Merkmale in German) is the set of attributes which forms the intension of the concept and I is a binary relation connecting G and M. We use the notation gIm (i.e. $(g,m \in I)$) which is read "the object g has the attribute m ". In the crosstable the rows are objects and the columns are attributes. An X indicates that a particular object has the corresponding attribute. Using the notion of a galois connection, formal concepts are found by determining the set of attributes shared by a set of objects or conversely the set of objects which share a set of attributes. Formally, a formal concept of the context (G,M,I) is defined to be a pair (A,B) with $A \subseteq G$, $B \subseteq M$, $A = \{g \in G | gIm \text{ for all } M \in B\}$ and $B = \{m \in M | gIm \text{ for all } g \in A\}$; A and B are called the extent and intent of the concept (A,B), respectively. The subsumption relation \geq is used to find sub-superconcept relations and to draw a complete lattice. In the FCA approach the rows in a crosstable typically represent a case with the columns showing the attributes of those objects. In our usage of FCA we first remove the MCRDR exception structure by picking up the conditions of all parent rules and then treat each flat rule as an object and the rule conditions as attributes. The objects are annotated by the rule number and the conclusion. The formal context and concept lattice for the conclusion %MC002 - "Metabolic compensation.2" from the Blood Gases domain are shown in Figures 3 and 4, respectively. Each object represents a primitive concept and higher level concepts are found by taking the intersections of the primitive rules. For more discussion on FCA and how it has been added to MCRDR see Richards and Compton 1997. For discussion on how we use different views to select which rules to include in our formal context and thus to display in a lattice see Richards (2000b).

	Normal Blood PH	Low Blood BIC	Low Blood PCO2	1=1	High Blood PH	High Blood PCO2	Low Blood PH	High Blood BIC	Incr Blood PH	Decr Blood BIC	Curr Blood PH≤7.36
9-%MC002	X	X	X	X							
10-%MC002		X		X	X						
14-%MC002	X			X		X					
15-%MC002				X			X	X			
19-%MC002	X			X		X			X	X	
49-%MC002	X			X		X					X

Figure 3: A formal context for the MCRDR rules which conclude %MC002- "Metabolic compensation.2" in the blood gases domain.

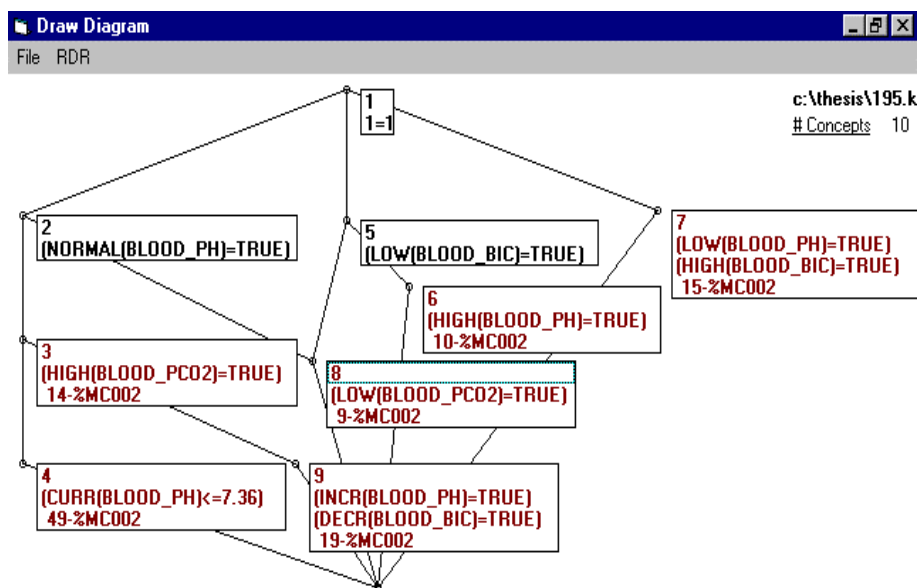


Figure 4: The Concept Lattice in MCRDR/FCA for the conclusion %MC002- "Metabolic compensation.2".

Each circle in figure 4 represents a concept. The attributes that belong to a concept are reached by ascending paths and objects are reached by descending paths. The rule conditions are the attributes. The objects are labelled using the rule number and the conclusion code. The conclusion codes have been made as meaningful as possible in the size limitation of five characters but a longer description could be substituted on the line diagram if desired for greater comprehension.

We have developed a framework based on MCRDR and FCA which allows multiple knowledge bases to be combined (Richards 2000a). The framework includes conflict detection, negotiation and resolution strategies. The benefit of having rules and cases is particularly useful in this situation. In the KA technique used in standard FCA the user is asked to offer a counterexample if they do not agree with the implications derived by the system. Coming up with such examples is often difficult. As can be seen in Figure 5, the case associated with concept number 8 has been popped up. This rule, which reads Expert C5 believes that IF [Silica=Very_High], [Colour=Light], [Grain_Size=Coarse] THEN the rock is Adamellite (%AD000). No other experts are in complete agreement with this expert. The ability to show the case gives the experts a concrete example to consider and saves the C5 expert from having to think up or remember an example.

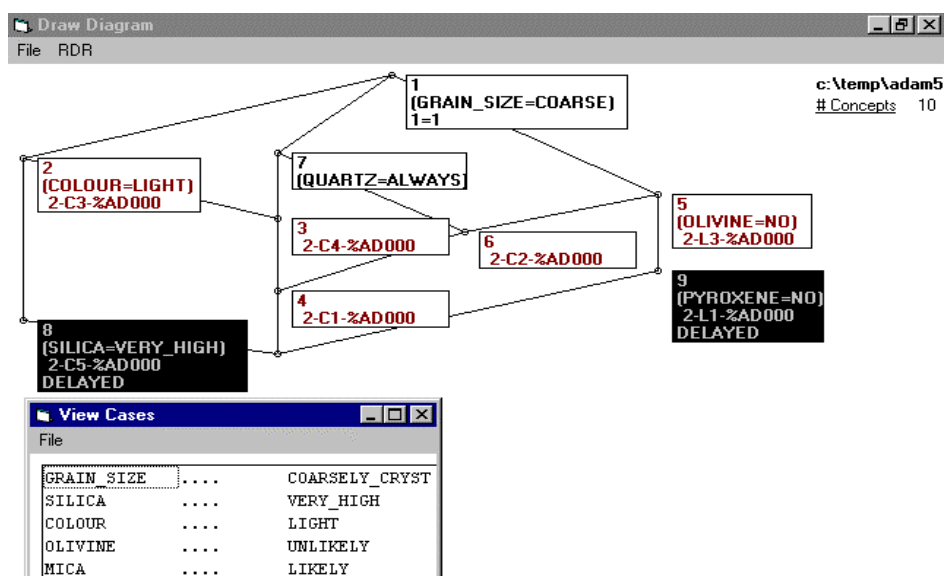


Figure 5: A line Diagram showing how a case can be popped up to assist in reconciliation of conflict between multiple sources of expertise.

FCA concept lattices are just one of the techniques we offer to assist in the transference of knowledge between individuals. We allow the user to perform knowledge acquisition in a critiquing mode. This allows the current user to see how the knowledge they propose to add fits in with the existing knowledge in the system. Another way to explore knowledge in an RDR KB involves the use of a nearest neighbour algorithm. Using this approach, the user can select a concept or a rule and then receive a list of other concepts or rules that are similar. A score is provided to allow the user to make comparison. Again the approach is aimed at providing the user with different ways of interacting with the knowledge. Rather than take a pedagogical approach where the system guides the learner, we do not assume to know better than the user what is relevant to them. By providing many different interfaces and representations of the knowledge we allow the user to develop and internalise their own models.

4. The Knowledge Management Challenge

Since tacit knowledge resides in an individual, effective KM must:

- minimise the loss/departure of individuals with large amounts of TK
- ensure that TK becomes externalised so that it can be transferred to others in the organisation (Van Daal, de Haas and Weggeman 1998)

We believe our approach can assist both of these goals. The approach that we offer is user-centred. Unlike many KBS approaches requiring a knowledge engineer to act as an intermediary, in our approach the user is a direct participant and has control and ownership of the knowledge. An RDR system is seen to be a decision support tool, not a replacement for human expertise. The Greek Oracle approach is not acceptable to experts (Miller and Masarie 1990) who are both suspicious of and insulted by systems that claim to be an Oracle. The ease with which knowledge can be captured and modified allows it to be tailored to suit individual preferences. Usability issues are part of the underlying philosophy of RDR (and FCA) and are key reasons for its success in the pathology domain. One reason knowledge is difficult to transfer is its *stickiness* (Johannessen, Olaisen and Olsen 2001; von Hippel 1994 in Ramaprasad and Rai 1996) which occurs because “a knowledge source may be reluctant to share crucial knowledge for fear of losing ownership, a position of privilege ..” (Szulanski 1996, p. 31). Indeed Sternberg in his seminal work on tacit knowledge measurement, refers to the attainment of tacit knowledge as typically occurring in conditions of ‘low environmental support’ (Sternberg *et.al.* 1995), in other words, the knowledge is not often willingly transferred from one individual to the other. Ownership and control by users and a situated view of knowledge have been major influences in the development of RDR (Richards and Compton 1998, Richards 2000c) which help to break down opposition to knowledge sharing.

In addressing the second point, the feasibility of our approach to capturing tacit knowledge that can be used by others has been demonstrated in numerous deployed systems. Unlike many KBS that have never left the research laboratory, we do not rely on the expert to state their knowledge. Current approaches which rely on the development of models in the form of ontologies ignore some fundamental shortcomings of models. Models are by their very nature approximations of the real thing and have been consistently found to vary between experts and even within an individual expert over time (Gaines and Shaw 1989). Validation of a conceptual model is very difficult and as a result, most approaches do not include any validation. A major problem with building and sharing models is that terms, their structure and meaning will differ. One of the reasons we were attracted to FCA was the use of an extensional and intensional definition. By describing an object by its set of features or conversely describing a feature by the objects that share that feature we have a little more information that can be used to decide if we are referring to the same concept. In our conflict resolution framework (Richards 2000a) we employ the four quadrant model of comparison of expertise developed by Gaines and Shaw (1989) which allows us to identify if two concepts are in a state of correspondence, consensus, conflict or contrast. This assists the user to determine which resolution operator to apply.

Knowledge exists in a social context and must be transferred from person to person. To achieve this, many recommend the use of apprenticeship or mentoring schemes within an organisation (Goldman 1990; Johannessen, Olaisen and Olsen 2001). As shown in Figure 1, in the RDR approach knowledge is transferred, but not directly person-to-person. The machine acts as the mediator and direct social interaction between the holder and receiver of the knowledge is not necessary. RDR formalises knowledge into a machine-readable representation which we restructure using FCA to uncover concepts not obvious or explicit in the original knowledge. Through exploration of the models and the various interfaces for knowledge exploration provided in our tool, the user is able to learn about the domain and develop their own models of how the knowledge fits together. We have demonstrated this in a couple of pilot studies. One study involved a domain beginner (level lower than novice) selecting parts of the knowledge to review and developing lattices for these views. The beginner then communicated with an expert from that domain to discuss what they had learnt. Four separate case studies were performed using a medical, agricultural, geological and chemical knowledge base. In every study it was found that the knowledge learnt was useful, with all experts stating that the beginner seemed to have deep understanding of the domain. In the medical domain, the expert’s response to one answer was “would you like a job as an anaesthetist?” In the other pilot study we tested how quickly someone could learn to read a concept lattice without prior experience and how well they could reason and form hypotheses for the domains shown in the lattices. Again the results were promising, showing that the ability to uncover abstractions and show relationships between concepts assisted in acquiring some deep knowledge about a domain.

The combined RDR/FCA approach is novel within the KBS community. There are other approaches which emphasise the role of the user, (e.g. the Protégé family of tools (Grosso *et al* 1994)) or which do not ask the expert to describe their knowledge but allows the knowledge to emerge through various interactions (e.g. tools based on personal construct psychology (Shaw and Gaines 1994)). However, in the first case there is still reliance on the user to define the knowledge models up front. In the second case, the techniques are not incremental in that user must consider the whole domain and specify the context at the start. The RDR approach is incremental and the context evolves as new cases are seen.

Knowledge management and explicit knowledge

Within the KM literature the closest work we have found to our own is a technique by Noh *et al.* (2000). That work is also case-based and works with cognitive maps which are similar in some respects to the FCA concept lattice. However, just as we have found in our comparisons with other work in the KBS area, Noh *et al.*'s approach begins with a formalisation phase in which the user is required to develop a cognitive map. The cognitive maps are stored in a case base. Given the shortcomings of models given above and since it is acknowledged by many that experts have trouble articulating their knowledge we have some reservations with starting with formalisation by the user. We also have a reservation regarding the cognitive maps themselves based on our experience into causal modelling which found that getting experts to formalise causal knowledge was extremely difficult since this knowledge was often unknown. A better approach was to automatically generate possible causal links and allow the user to review and revise these (Lee and Compton, 1995). Kolodner (1993) suggests the use of cases as the starting point in domains where causal models are not well understood. However, in Noh *et al.*'s approach causal knowledge must be acquired first from which cases are developed. From Kolodner's remarks we could conclude that the knowledge being captured is actually explicit rather than tacit knowledge. Following the formalisation phase in Noh *et al.*'s approach is the reuse phase where *fitting* and *garbage* ratios are used to retrieve appropriate cognitive maps from the case base which are adapted to fit the new situation. Indexing, retrieval and adaptation of cases are not simple tasks. To overcome these difficulties, the RDR approach uses rules specified by the expert in the course of problem solving as the indexes to our case-base. The final phase of Noh *et al.*'s approach is problem solving where the adapted cognitive model is applied to the new problem and then stored in the case base. A comparison with our process reveals that we begin rather than end with problem solving and that our formalisation phase is handled by the system rather than the user.

5. Conclusion

The identification and transfer of TK is of paramount concern in today's knowledge economy. Many have recognised this but few have been able to offer solutions. We offer a solution in this paper based on our experience with KBS. The solution is fundamentally different to other research we have encountered, most of which assumes the solution to be codification or communication of the knowledge to others. The solution appears to be the problem. If we consider TK to be that knowledge which is difficult to communicate as information, entrained in action (practice) (Johannessen, Olaisen and Olsen 2001) and linked to concrete contexts (Rolf 1995 in Johannessen, Olaisen and Olsen 2001) then our knowledge acquisition approach which is based on concrete cases and the expert practising their expertise without the need to communicate it directly to another human seems to fit nicely. It does not make sense to develop a process that relies on human codification of something we consider to be either inarticulable or difficult to articulate. With RDR we capture knowledge in action directly from the human expert. The high degree of participation, ownership and control afforded by the RDR technique together with the simplicity of the approach encourages user satisfaction and willingness to use the system. Others within the organisation can utilise that knowledge by executing the rules or they can develop and compare models to internalise that knowledge using such techniques as FCA. In either case, the result is that the knowledge stays in the organisation after the individual expert retires or leaves.

Bibliography

Busch, P. and Richards, D., (2000) Measurement of articulable Tacit Knowledge using Formal Concept Analysis *Australasian Conference on Information Systems (ACIS2000)* Queensland University of Technology December 6th – 8th

- Chandrasekaran, B. and Johnson, T. (1993) Generic Tasks and Task Structures In David, J.M., Krivine, J.-P. and Simmons, R. (eds.), *Second Generation Expert Systems* Springer, Berlin, 232-272.
- Clancey, W.J., (1992) Model Construction Operators *Artificial Intelligence* 53:1-115.
- Clancey, W.J., (1993) Situated Action: A Neurological Interpretation Response to Vera and Simon *Cognitive Science*, 17: 87-116.
- 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, 247-291.
- Compton, P. and Jansen, R., (1990) A Philosophical Basis for Knowledge Acquisition. *Knowledge Acquisition* 2:241-257.
- Edwards, G., Compton, P., Malor, R., Srinivasan, A. and Lazarus, L. (1993) PEIRS: a Pathologist Maintained Expert System for the Interpretation of Chemical Pathology Reports *Pathology* 25: 27-34.
- 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.
- Goldman, G., (1990) The tacit dimension of clinical judgement *The Yale Journal of Biology and Medicine* 63(1) January/February :47 – 61
- Grant, E. and Gregory, M., (1997) Tacit knowledge, the life cycle and international manufacturing transfer *Technology Analysis & Strategic Management* 9(2) :149 – 161
- Grosso, W. E, Eriksson, H., Ferguson, R.W., Gennari, H. Tu, S.W. and Musen, M. (1999) Knowledge Modelling at the Millenium (The Design and Evolution of Protégé-2000) *In Proceedings of the 12th Workshop on Knowledge Acquisition, Modeling and Management (KAW'99)*, 16-21 October, 1999, Banff.
- Gruber, T. R., (1993) *Toward Principles for the Design of Ontologies Used for Knowledge Sharing* Technical Report Knowledge Systems Laboratory, Stanford University.
- Guha, T.V., and Lenat, D.B. (1990) CYC:A Mid-Term Report *AI Magazine* 11(3):32-59.
- Howells, J., (1995) *Tacit knowledge and technology transfer* Working paper No. 16 ESRC Centre for Business Research and Judge Institute of Management studies, University of Cambridge U.K. September
- Johannessen, J., Olaisen, J. and Olsen, B., (2001) Mismanagement of tacit knowledge: The importance of tacit knowledge, the danger of information technology and what to do about it *International journal of information management* Vol 21 :3 – 20
- Kang, B. (1996) *Validating Knowledge Acquisition: Multiple Classification Ripple Down Rules* PhD Thesis, School of Computer Science and Engineering, University of NSW, Australia.
- 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.
- Kolodner, Janet (1993) *Case-Based Reasoning* Morgan Kaufman Publishers Inc., San Mateo, CA.
- Lee, M. and Compton, P. (1995) From Heuristic Knowledge to Causal In X. Yao (ed), *Explanations Proceedings of Eighth Australian Joint Conference on Artificial Intelligence AI'95*, 13-17 November 1995, Canberra, World Scientific, 83-90.
- McCarthy, J. (1991) *Notes on Formalizing Context* Technical Report. Computer Science Department, Stanford University.
- Miller, R.A. and Masarie, F.E. (1990) The Demise of the Greek Oracle for Medical Diagnosis Systems, *Meth Inform Med* 1990 (29):1-2.
- Newell, A. (1982) The Knowledge Level *Artificial Intelligence* 18:87-127.
- Newell, A. and Simon, H.A. (1976) Computer Science as Empirical Inquiry: Symbols and Search, *Communications of the ACM*, 19(3), 113-126.
- Noh, J., Lee, K., Kim, J., Lee, J. and Kim, S., (2000) A case – based reasoning approach to cognitive map – driven tacit knowledge management *Expert systems with applications* Vol. 19 :249 – 259
- Nonaka, I., Takeuchi, H. and Umemoto K., (1996) A theory of organisational knowledge creation *International Journal of Technology Management* 11(7/8) :833 – 845
- Norman, D.A. (1988) *Psychology of Everyday Things* Basic Books, New York.
- Patil, R. S., Fikes, R. E., Patel-Schneider, P. F., McKay, D., Finin, T., Gruber, T. R. and Neches, R., (1992) The DARPA Knowledge Sharing Effort: Progress Report In C. Rich, B. Nebel and Swartout, W., *Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference* Cambridge, MA, Morgan Kaufman.
- Pylyshyn, Z., (1981) The imagery debate: Analogue media versus tacit knowledge In *Readings in Cognitive Science: A perspective from Psychology and Artificial Intelligence* (eds. Collins, A., Smith, E.,) Chapter 6.5 Morgan Kaufman San Mateo California U.S.A. :600 – 614
- Raghuram, S., (1996) Knowledge creation in the telework context *International Journal of Technology Management* 11(7/8) :859 – 870

- Ramaprasad, A. and Rai, A. (1996) Envisioning management of information *Omega: International Journal of Management Science* 24(2) April :179 – 193
- Reber, A., (1993) *Implicit learning and tacit knowledge: An essay on the cognitive unconscious* Oxford Psychology Series No. 19 New York Oxford Oxford University Press Clarendon Press
- Reed, S. and Hock, H., (1983) Tacit knowledge and the effect of pattern recognition on mental
- Richards, D. (2000a) Reconciling Conflicting Sources of Expertise: A Framework and Illustration In *Proceedings of the 6th Pacific Knowledge Acquisition Workshop* P. Compton, A. Hoffmann, H. Motoda and T. Yamaguchi (eds.) Sydney December 11-13,2000, 275 - 296.
- Richards, D. (2000b) The Visualisation of Multiple Views to Support Knowledge Reuse *Intelligent Information Processing Conference (IIP'2000)* In conjunction with the 16th IFIP World Computer Congress WCC2000, August 21-25, 2000, Beijing, China.
- Richards, D. (2000c) The Reuse of Knowledge: A User-Centered Approach, *International Journal of Human Computer Studies* 52 (3) March 2000: 553-579.
- Richards, D and Compton, P., (1997) Combining Formal Concept Analysis and Ripple Down Rules to Support the Reuse of Knowledge *Proceedings Software Engineering Knowledge Engineering SEKE'97*, Madrid 18-20 June 1997, 177-184.
- Richards, D and Compton, P. (1999) An Alternative Verification and Validation Technique for an Alternative Knowledge Representation and Acquisition Technique, *Journal of Knowledge Based Systems* 12 (1999) 55-73.
- Richards, D and Compton, P. (1998) Taking Up the Situated Cognition Challenge with Ripple Down Rules *International Journal of Human Computer Studies, Special Issue on Situated Cognition.* (49) 895-926.
- Schon, D.A. (1987) *Educating the Reflective Practitioner* Jossey-Bass, San Francisco, CA.
- Shaw, M.L.G. and Gaines, B.R.. (1994) Personal Construct Psychology Foundations for Knowledge Acquisition and In Steels, L. Schreiber, G. and Van de Velde, W. (eds) *A Future for Knowledge Acquisition Proceedings of the 8th European Knowledge Acquisition Workshop, EKAW'94*, Lecture Notes in Artificial Intelligence 867, Springer Verlag, Berlin, 256-276.
- Sternberg, R., (1999) Epilogue – What do we know about tacit knowledge?: Making the tacit become explicit In *Tacit Knowledge in Professional Practice: Researcher and Practitioner Perspectives* (Eds. Sternberg, R., Horvath, J.) Lawrence Erlbaum and Associates Mahwah New Jersey U.S.A. :231 – 236
- Sternberg, R., Wagner, R., Williams, W. and Horvath, J., (1995) Testing common sense *American psychologist* 50(11) November :912 – 927
- Szulanski, G. (1996) Exploring internal stickiness: Impediments to the transfer of best practice within the firm. *Strategic Management Journal*, 17 (winter special issue), 27-43.
- van Daal, B., de Haas, M. and Weggeman, M., (1998) The knowledge matrix: A participatory method for individual knowledge gap determination *Knowledge and process management* 5(4) :255 – 263
- 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. (1992) Concept Lattices and Conceptual Knowledge Systems *Computers Math. Applic.* (23) 6-9: 493-515.
- Wille, R. (1996) Conceptual Structures of Multicontexts In P. Eklund, G. Ellis and G. Mann (eds.) *Conceptual Structures: Knowledge Representation as Interlingua* *Proceedings of the 4th International Conference on Conceptual Structures*, LNAI 987, Springer. Verlag, 23-39.