Chapter 9

9 Thesis Conclusion

In this chapter we review the work described in this thesis, its contribution, the limitations and possible future work.

9.1 A Summary of the Work Presented

The work reported in this thesis is motivated by the belief that we need to focus more on users' needs and cater for the various decision situations in which users will find themselves. To build individual systems that cater for all the activities that may be needed is not feasible or desirable. The problems associated with capturing knowledge are well known and the ability to capture knowledge once and access and manipulate the knowledge in multiple ways adds value to the original knowledge and offers all the benefits associated with the reuse of resources. Thus the problem becomes one of knowledge reuse. The research question pursued in this thesis was "could knowledge captured for one activity be reused to support a wide range of alternative activities allowing the user to answer different types of questions" ? Further, this goal was to be sought in a situated cognition, dynamic knowledge framework.

The scope of this work has been broad rather than deep as the number of *reuse-activities* considered was extensive and the goal of this work has been to see how many different activities could be supported using the same knowledge rather than focusing on a good or best solution to any one particular activity. The use of RDR has allowed greater attention to be paid to the reflective activities that require some understanding and explanation of the knowledge since in RDR the reflexive activities of KA, inferencing and maintenance were already well established within the one system. The strengths and shortcomings of RDR were presented in Chapter Three together with a comparison of RDR with other similar techniques, notably personal construct psychology and case-based reasoning. Also presented in Chapter Three was some work that was done on rule-based reasoning which resulted in the development of a nearest-neighbour algorithm to determine the closeness of concepts within the knowledge base. The limitations of RDR seen as relevant to *activity-reuse* were repetition, lack of higher-level models and the handling of multiple sources of expertise and have been addressed in Chapters Four, Five and Six.

Chapter Four addressed the problem of repetition of knowledge in RDR KBS. Some experiments were described that sought to compact RDR KBS by removing redundancy and repetition. It was expected that the cleaned-up knowledge would facilitate reuse of the knowledge and particularly improve the quality of the explanations provided. The experiments reported had mixed results and the benefits of removing repetition were not clear. A major limitation of the experiments was that they only considered single-classification RDR and it is conjectured that the ambiguity of the false branches was a major cause for the poor results. It also appears that the Garvan data used in the study was not statistically distributed which is inappropriate for a statistical method such as Induct. A proposal for the removal of repetition from MCRDR KBS was offered but further exploration of that proposal was left to future work. The fourth chapter also offered an extension to the RDR approach to verification of MCRDR KBS through a novel usage of rough set theory. The use of RST for verification was also offered as a general approach applicable to any KR that could be converted to decision table format.

Chapter Five described a solution to the lack of higher level modelling available in RDR which was important for many of the reflective activities. This work, together with the development of the nearest neighbour algorithm described in Chapter Three, allowed the relationships between concepts to be uncovered. Through the incorporation of FCA into MCRDR it was possible to develop an abstraction hierarchy from the performance rules. This allowed abstractions and the underlying structure of the rules, which were not previously visible, to be found.

Chapter Six reported on an extension of the approach in Chapter Five to support requirements engineering. This RE work addressed a shortcoming of RDR not considered in any depth in other RDR research which was "how to handle multiple sources of expertise ?". Also, this work offered an additional *activity-reuse* of the knowledge which was "how does my knowledge compare to knowledge from one or more other sources ?" and "how can differences between the viewpoints be reconciled ?" Chapter Six demonstrated the extendability and possible widespread application of the use of FCA for KBS analysis by building terminological KBS from assertional KBS. The RE approach developed has also proposed some new strategies for the maintenance of MCRDR KBS based on the identified changes in the terminological KBS. As part of these strategies, the paths that ended with stopping rules were not included in the terminological KBS but if the action required a modification rather than a deletion, such as removing one condition, a new rule with the correct conclusion would be added at the top of the MCRDR tree with all conditions bar the dropped one.

The aim of Chapter Seven was to give the reader a feel for the total product of the work in the prior chapters and to demonstrate that the original goal of this thesis, to build a multi-purpose system that uses and reuses knowledge for different activities, had been achieved. The chapter was not designed to be a comprehensive guide to each activity or to be a definitive description of the system but to offer a view of how a number of activities could be included in one system without the need for changes to the knowledge or the inference strategy. The chapter embodies what Stelzner and Williams (1988) had envisaged when they said that KBS should be structured so that changing the usage of the knowledge only required changes to the user interface. The original MCRDR system could support a number of activities already and FCA provided the link between what was already available and those activities such as critiquing and teaching that needed abstractions and structure to be made explicit.

Chapters Three, Four, Five, Six and Seven included sections which evaluated the work reported in the chapter. Chapter Eight concentrates on an evaluation of the FCA line diagram as its value is pivotal to much of this thesis and from discussion with the developers of FCA in Darmstadt such an evaluation had not previously been performed.

In this final chapter, the contribution of this thesis is considered including not only what has been delivered but what may be possible.

9.2 The Contribution of this Thesis to KBS Research

This thesis has made a number of contributions. Some of them relate directly to the original goal and some have been outgrowths of the work. Two of these outgrowths are the use of rough set theory for verification and the development of a requirements engineering framework, both of which apply beyond RDR KBS. While both of these explorations were detours from the original goal of reusing knowledge for different activities and have therefore not been pursued as far as possible, the techniques described in this thesis are comprehensive and supported by implemented systems and examples. These detours were natural consequences from the investigation of the use of RST for finding relationships between concepts and the use of MCRDR/FCA for conflict resolution and are evidence of the usefulness of the approaches developed beyond their intended use in addressing the concerns of this thesis.

The main contribution of this thesis that is seen as more directly applicable to the research question is the incorporation of FCA into MCRDR to build abstraction hierarchies from the primitive rules. This contribution has been significant because it has demonstrated that it is possible to start with a performance system and derive an explanation system. This finding is in stark contrast to the thrust of most KBS research which has concentrated on complex modelling of the domain and problem solving knowledge before KA could even commence. It has been argued in this thesis that the imprecise nature of models and the situated nature of knowledge has resulted in approaches that have done little to alleviate the KA bottleneck. While the general problem solving method and ontological approaches based on the knowledge level are superior to the earlier 'transfer-of-knowledge' approaches, they still seem to indicate an Platonic view that knowledge is an artifact and that if we can only model it correctly the rest will fall into place.

We are beginning to a see a change in this view. At the recent EKAW'97 in Barcelona different research groups were starting to view different problem solving methods as variations of a limited number of search strategies. This change is welcome. The focus on building the right model isolated the expert/user even further from the system and has resulted in systems based on the KE's model rather than the experts. Less focus on these models and a greater focus on the data and the development of performance systems seems more realistic and practical. The RDR philosophy is itself undergoing change as it presents itself as a family of problem solving methods where the PSM is weak (that is, it is a general-purpose strategy) and relies on enough and properly-structured data to be thrown at it to support tasks of classification and construction. It appears that both camps are moving to more middle ground.

This thesis has argued that while explanation systems are difficult to capture and even harder to validate they are useful and necessary for many of the activities that a user will need to perform. This thesis offers a system that includes both performance and explanation knowledge by starting with an assertional KBS, that can be captured and validated with minimal effort, and later deriving the terminological system. FCA is used to provide automated reflection on the primitive concepts in the performance system by uncovering some of the higher level concepts that experts often use but have trouble articulating. The models built by FCA can be presented to the expert and used as a further source of validation and to assist the expert with KA where reflective action is necessary. This thesis has also argued that by making explicit the relationships between concepts in the KBS and by structuring them into an abstraction hierarchy the user is better able to use the knowledge in the KBS for explanation, critiquing and learning purposes.

Another, but less explored, contribution is the use of the nearest neighbour algorithm offered in the MCRDR/FCA system which provides a way of determining the closeness of concepts. Its reliance on comparison of strings within a rule is a limitation but this limitation appears unavoidable unless some other type of symbol or number is used to encode the relationships. In MCRDR/FCA the nearest neighbour algorithm is used to determine which other rule pathways are closest by providing a proximity score between 0-1. This is seen to be useful for providing KA as a reflective critiquing activity and could also be used in learning more about a domain such as in the example in Chapter Three where the higher the score the closer the animal listed was to the man concept selected. Some potential uses of this algorithm, which depending on the situation may require modification, are suggested in the following section. While the nearest neighbour algorithm is part of the MCRDR/FCA toolkit there is no claim that this algorithm is original or optimal. This would require more thorough investigation of other similar algorithms and more support from theoretical semantic analysis. The algorithm was developed from a brief review of cluster analysis research and appeared to be suitable for the task for which it was designed. For example, a comparison between the algorithm and other work by Mehrotra (1995a and 1995b), who has developed a multi-viewpoint clustering analysis methodology, revealed that the algorithm developed in Section 3.3.2.1 provided most of the features offered in the clustering analysis methodology.

Of equal importance is the contribution of this thesis in bringing together a number of theories and techniques. In many cases comparisons and evaluations have been offered in this thesis. The theories and techniques include: Induct, Rough Set Theory, Personal Construct Psychology, Case-Based Reasoning, Semantic Networks and Formal Concept Analysis. The closeness of many of the techniques was interesting to the author particularly their use of differences and/or similarities to find concepts. I was intrigued whether consideration of differences or similarities was the stronger approach or what were the various merits of each approach. This question was clearly and simply answered following a keynote speech by Dedre Gentner at the 4th Australasian Cognitive Science Conference who referred me to a number of papers and pointed out that differences only become important once some similarity has been established. When I considered RDR, PCP, RST or FCA I could see that all approaches either explicitly or implicitly accepted this theory. Only when a case is similar is it and a difference list shown to the user in RDR. In PCP, triadic elicitation requires the user to state what two elements are alike and what third element is different. In RST sets of objects which share the same features are found first and then used to determine differences. An FCA formal context is essentially a grouping of related objects which are being differentiated by the attributes that describe them.

It is noteworthy that no new theory has been developed in this thesis. It appeared to me that there are many ways to solve a particular problem and also that many theories are very similar. Pawlak (personal communication) has commented on the similarity between the use of set theory in RST and FCA with the main difference being what they were trying to achieve. In this thesis I have tended to consider existing theory and to determine how it can be used or adapted to solve the problem at hand. I think that as a community we should work together borrowing from one another rather than reinventing the wheel if one was already available. As a community we suffer when we are inundated with a myriad of alternative approaches with no clear way of determining which one is the best. We also have a chance at achieving some standards if we stabilise the approaches we are using.

Many of the lessons learnt in the MYCIN experiments are confirmed in the work performed as part of this thesis. These lessons included the need for abstracted and more knowledge relevant to the particular activity to be performed. What differentiates the work in this thesis from Clancey's work in GUIDON is that the relationships between concepts and the higher level abstractions, which is some of the extra knowledge that is needed, can be derived automatically using FCA. In Clancey's approach additional knowledge was coded into rules which had different roles, such as the T-rules used to guide the tutoring process. This difference constitutes a substantial benefit because of the reduction in effort required and, as already noted, the capture of performance knowledge seems easier and more reliable than the capture of explanation or terminological knowledge. There is also high fidelity between the performance and explanation systems and ontological commitment between the KBS and FCA built ontology since one is based directly on the other.

This section would not be complete without some analysis of whether a multi-purpose system that allowed the user to reuse knowledge in a wide range of activities had been achieved. The MCRDR/FCA system supports the activities of KA, maintenance, inferencing, critiquing, explanation, tutoring, 'what-if' analysis and hypothesis testing. While it is not claimed that these activities are fully functional or ideally presented, the key concern was whether the knowledge content and structure were adequately robust to allow reuse for different activities. Initial examination of RDR showed that while many activities could be supported to some extent, it was necessary to find a method for finding concepts, both primitive and abstracted, and the relationships between them to provide a full range of activities. The use of FCA, and to a lesser extent the nearest neighbour algorithm, have addressed this need.

The activities offered in MCRDR/FCA are designed to exhibit transparency of interaction (Winograd and Flores 1986) or *readiness-to-hand* (Heidegger 1962), meaning that the system should appear as an extension to the human assisting them in the task they are performing. The user may switch between activities without a conscious change, being driven by their circumstances to ask a different type of question of the system. The activities supported by MCRDR/FCA covers reflexive and reflective activities because humans will at different times behave in these different ways.

9.3 Limitations and Future Work

While this thesis provides a proof-of-concept that one can use and reuse knowledge for multiple purposes and that we can start with a performance system and derive an explanation system there is still much more work to be done particularly in the area of HCI. The user evaluation of the MCRDR/FCA line diagram performed in Chapter Eight was favourable but inconclusive. To a large extent the decision of what should be included in a formal context, the use of the concepts derived, the use of the score provided by the nearest neighbour algorithm and the required features of the line diagrams need to be user and application driven. Without greater user involvement the system developed can not be more than a prototype that can act as a tool for communication of ideas. Nevertheless, some improvements are suggested. Some of the features that would greatly increase the utility of MCRDR/FCA would be the ability to scroll off the screen to other parts of the diagram, to zoom in and out at specified nodes or the ability to drop or add nodes as desired. This would particularly enhance the activities that require explanations such as ‘what-if’ or scenario analysis. A simple improvement would be to display the short description of the conclusion instead of the conclusion code as the object name. The code makes comprehension of the meaning of implications and ‘isa’ relationships difficult. As suggested in Section 8.1.1.2, an interesting extension to the MCRDR/FCA tool would be to add a run-time feature that selected which parts of the KBS to add to the formal context based on the current case or a user-response. It is also desirable to allow a split screen so that more than one diagram could be viewed at a time for comparison, although this would require larger screens in many cases. A comparison of the layout algorithm to other algorithmic and/or deductive approaches (Eades 1996) should be made with the possible facility for the user to select from a range of layout methods. Currently the layout algorithm works fairly well as a first approximation but depending on the size and shape of the labels, nodes may need to be moved to avoid overlapping.

More theory-related considerations of the limitations of the approach offered is the simple way that MCRDR KBS, or other rule-bases for that matter, are converted into a formal context. As described in Chapter Five, the formulation of concept lattices from many-valued contexts requires their interpretation into a binary formal context. The interpretation of a rule condition as an attribute in the crosstable has proven straightforward and sufficed for the purposes of the prototype developed. However, a rule that should be part of the context for a selected focus of attention may be missed if the condition does not match on a conclusion or attribute already selected. The use of different conceptual scales may provide a solution and needs further investigation. Perhaps the distance-weighted nearest neighbour algorithm described in Chapter Three could be used to assign a score of relative proximity to indicate which clauses are related and to what extent. This score could then be used to determine which rules should be added to a context. Another option is to generate concepts for the whole KBS and use the FCA concept hierarchy to determine the distance between nodes. Those within a specified distance would be added to a second formal context from which a smaller set of concepts could be derived.

The removal of repetition from MCRDR is still an open question and one that should be pursued. In particular the strategy for the handling of stopping rules needs to be investigated further. From the work that I have done using FCA to draw line diagrams of MCRDR KBS, the stopping rules were often hard to interpret or explain. In some cases it appeared that a stopping rule meant if a case has certain additional conditions then the previous conclusion should no longer be given otherwise the old conclusion was still valid. In other cases the stopping rule may mean that the previous rule, and the entire pathway associated with it, should be ignored. This is a similar problem to the ambiguity of the false branches. With the false branches you do not know if the case fails because it was a 'don't care' situation or because the condition should be false. For the purposes of explanation it is important to know the purpose of the stopping rule and what the appropriate handling should be. If the previous rule is to be stopped then the pathway should be ignored. If the previous rule still holds but not under certain exceptions then it should be retained together with its exceptions. Recent discussions with the developer of MCRDR indicate similar reservations. It would have been good to address in this thesis the shortcomings of the stopping rule strategy and the MCRDR repetition problem. However, the MCRDR representation needs time to mature and more MCRDR KBS to be built before these issues can be adequately addressed.

In Chapter Two, ontologies and some of their limitations were presented. The main problems were that they are time consuming to capture, they are difficult to validate and they can be difficult to change. The abstraction hierarchy built by MCRDR/FCA can be seen as an ontology as was demonstrated in the examples in Chapter Eight which used RDR to build taxonomies. However, in the case of MCRDR/FCA none of the limitations mentioned exist because the ontology is built retrospectively based on the performance knowledge that is easier to capture and validate. This feature, however, means that the descriptiveness of the ontology derived is limited by the assertions and can be considered as a first step. A descriptive model (T-box) may require deeper knowledge to be captured and represented. As was mentioned in section 7.3.2, causal knowledge is an example of deeper knowledge that needs to be acquired but this is still consistent with the approach as the abstraction hierarchy acts as a causal model based on the causal assertions.

There may be merit in the addition of higher level concepts as new rules but since this information can be derived it would add to the incidence of repetition and redundancy in the KBS. Of greater concern is how the higher level knowledge would be maintained since changes to the primitive rules can result in a change in the higher level concept. Maintenance of abstractions is not a concern when only the primitive rules are kept in the assertional knowledge base and the terminological knowledge base is regenerated each time it is needed or when changes are made to assertions.

To support the reuse of abstractions, it is possible with MCRDR/FCA to attach a name to a higher level concept which is stored in a file and whenever a concept is found which shares the same set of attributes, the concept name is shown on the line diagram and can also be accessed via the concept matrix screen. In this way higher level concepts can be identified and reused. This is a similar goal to that sought by NRDR which was described in Section 3.1.3. However, as noted in the previous paragraph, maintenance of abstractions can pose a problem as is the case for NRDR. Taking usage of abstractions a step further, Compton (personal communication) has suggested that these intermediate and high level concepts could be used to support reasoning at different levels. As commonly advocated in the literature, this would be particularly useful for explanation purposes but could also be useful in providing a recommendation depending on the type of information the user was after. The main problem then becomes how to decide what level should be offered. To avoid the problems associated with user-modelling, the simplest, and quite possibly the best, alternative is to let the user specify the level of reasoning and to allow the user to change the level.

Menzies (personal communication) has proposed two further extensions of the use of MCRDR/FCA. The first one he calls Ripple-Down Ontologies. The idea is that we take the higher level concepts found in a T-box derived from an MCRDR/FCA A-box and see if we can reuse these in a new application that appears to have similar concepts. For example if the lower concept “cat” in the first KBS has the features: suckles young and warm-blooded which results in the higher concept “mammal” we may find a lower concept “dog” in a second KBS that has these two attributes which prompts us to assume that we are dealing with a similarly structured domain. We could then copy over the structure developed in the first KBS and modify the copy incrementally to suit the second application. This would result in a bottom-up initial approach followed by a subsequent top-down approach to system development. This top-down approach, like that used by Martinez-Bejar, Benjamins and Martin-Rubio (1997) can be used to assist the user with the population and validation of the KBS and could also be a way of combining different sources of expertise. The first step of identifying similar concepts can be easily performed through the use of the concept name file which was mentioned above. Bringing over all of the higher level concepts (not just the ones that match) has not been tried and would require the objects (the rules) from the first KBS to be added as there must be at least one example of every higher level concept. This will introduce lower level concepts that may not be of relevance to the second application.

Another interesting suggestion by Menzies is the use of MCRDR/FCA for evaluating ontologies. Menzies points out that there is an assumption within ontologies that reusable terms exist. To test this we can give the same cases (which ensures consistent terminology), in the same order, to different experts and compare the RDR rules and FCA T-boxes that they build. If the T-boxes developed are different and cannot be resolved using requirements engineering we can reject the use of ontologies. This would be an interesting exercise which has not been performed due to the unavailability of multiple experts in the same domain, to conduct such an experiment.

A number of uses of the MCRDR/FCA tool have been suggested by Compton (personal communication). One is the use of the abstraction hierarchy to determine which conclusions are similar or the same which could then be combined to reduce the complexity of the KBS. Alternatively the nearest neighbour algorithm could be used on the rule conditions and/or the conclusion descriptions to find how close the conclusion strings are to one another. The use of MCRDR/FCA to reduce the KBS complexity seems feasible but has not been performed due the size limitations of Visual Basic in the handling of files, grids and strings. There were two KBS available, one with 600 the other with 1300 rules that could have possibly been reorganised for this exercise but the limitation of Visual Basic meant the current tool was unsuitable. A reimplementation in another language was not desirable at this stage.

Another problem being faced by a pathology MCRDR KBS currently under construction (3,000 rules to date) is the situation where the same conclusion is being reached but through different pathways. This is satisfactory for inferencing as duplicate conclusions are edited out but for the purposes of explanation it is unclear which path should be reported and how to explain the different reasons. Multiple pathways also make KA difficult as there are numerous situations to consider. The MCRDR/FCA hierarchy could provide a way of ordering the rule pathways to reveal what were the main concepts and what variations existed.

It is apparent that the MCRDR/FCA tool has opened up the way to addressing a number of problems concerning the use and understanding of knowledge in a KBS. These issues have not been addressed in this thesis and those mentioned are only a few examples of the potential value of such a tool. The realisation of this potential is left to future work.

9.4 Final Words

As described, many KBS approaches assume the domain must be fully described before we can begin to acquire knowledge and seem to take the view that "there is no perception without prior learning" (Rosenfield 1988, p.7). However, the focus of KBS should be on systems that support change and evolution of the knowledge since new perceptualisations of categories are continually reconstructed. This also means that it is not necessary or feasible to try to encode all the knowledge before it can be put into use as most approaches to KBS development require (Kang, Gambetta and Compton 1996). Systems can be put into routine use with a minimal set of rules and developed online, as was the case with PEIRS. It also means that rather than using KBS methods that rely on identification of the task structure or problem solving method/s before the appropriate method can be applied it is desirable to have a KA and representation technique that allows a system to be built incrementally and "gradually evolve into whatever type of ... system is best suited to the problem as new cases are seen" Kang, Gambetta and Compton 1996, p.267, which was first suggested in Compton et al (1993).

Just as RDR stresses the incremental nature of KBS development, MCRDR/FCA systems are not expected to be complete but able to evolve. This evolution allows new activities to be added according to user requirements. The key concern is whether the MCRDR/FCA framework is sufficiently robust so that new activities only require user interface enhancements or more knowledge. As was found by Lee and Compton (1995), causal modelling required additional knowledge but it was not necessary to change the structure of the knowledge or the way that inferencing was performed. Causal modelling was possible due to the causal positive or negative relationships between nodes. These relationships could be derived automatically or captured manually. The recent work (Ramadan et al 1997a) in configuring for ion chromatography supports the reuse of the MCRDR algorithm for classification and configuration tasks. MCRDR can handle well the reflexive activities of KA, maintenance and inferencing. By adding FCA tools to MCRDR, MCRDR/FCA is also able to support reflective activities such as explanation, critiquing and 'what-if' analysis. Thus, MCRDR/FCA supports the reuse of knowledge for multiple activities. In the case of PSM and activity reuse, data is the key. It is anticipated that if sufficient data of the right type is captured, RDR, including MCRDR/FCA will be able to handle whatever task or activity is required. This is a big claim and one that is not substantiated in this thesis. But this is the direction in which RDR research is heading and MCRDR/FCA also pushes RDR in this direction.

The system described is an attempt to build a system that does not force the system developer's prejudices onto the user by restricting interaction to a particular type of activity. Accepting the limitations and developer-imposed restrictions in the user interface, the idea is that the user can ask their own questions rather than having to reframe their question into the type of activity that is offered by the system. By offering a wide range of ways of utilising the knowledge the user has the flexibility to interact in the way that best suits their current situation.

The work and tool described in this thesis were aimed at building a system that would allow the user to view, explore, analyse, maintain, manipulate and consult the knowledge in a knowledge based system. While the tool itself has many limitations, the key aspect to be emphasized is the freedom on the part of the user to decide how to use this tool and is well expressed in the following idea of Winograd and Flores:

> *"the use of the tool [referring to the computer] shapes the potential for what those actions are and how they are understood....its power does not lie in a single purpose."*
> *Winograd and Flores 1986, p.170."*

It is also believed that the system described embodies the vision that:

> *"The most successful designs are not those that try to fully model the domain in which they operate, but those that are 'in alignment' with the fundamental structure of that domain, and that allow for modification and evolution to generate a new structural coupling." Winograd and Flores 1986, p.53.*

RDR allows the user to capture a simple observable model of their world using attribute-value pairs and conclusions. By allowing the expert to retrospectively view and explore their underlying conceptual models using FCA, together with a situated system style as described above, it is expected that the system may be a vehicle for transformation of tradition (Winograd and Flores (1986, p.170). We all act according to a tradition based on a communal view. By helping the expert understand their domain we help them to understand, question and possibly change their tradition. Much of the failure of ES is due to the improper view of what they are and should be able to achieve. An ES must be seen as a tool and the emphasis must be on support not on replacement. In this thesis the emphasis has been on giving the user a flexible tool that the user controls by choosing how to use and reuse the knowledge according to their situation and personal preferences.