Reverse Engineering Ontologies from Performance Systems

Debbie Richards
Department of Computing
Division of Information and Communication Sciences
Macquarie University
Sydney, Australia
Email: richards@ics.mq.edu.au

Abstract

Considerable effort is associated with the development, validation and integration of ontologies. This paper suggests that an alternative, or possibly complementary approach, to engineering ontologies is to retrospectively and automatically discover them from existing data and knowledge sources in the organization and then to combine them if desired. The method offered assists in the identification of similar and different terms and includes strategies for developing a shared ontology. The approach uses a data analysis technique known as formal concept analysis to generate an ontology. The approach is particularly strong when used in conjunction with a rapid and incremental knowledge acquisition and representation technique, known as ripple-down rules. However, any data that can be converted into a crosstable (a binary decision table) can also use the approach. The ontological representation is not as rich as many others but we have found it useful for uncovering higher-level concepts and structure that were not explicit in the performance data. If richer models are required our approach may provide a quick way of developing a first draft and gaining initial ontological commitment.

1. An Overview

The alternative technique proposed for ontological engineering in this paper begins with rapid development of a performance system using Multiple Classification Ripple Down Rules (MCRDR) [14] followed by automatic generation of an ontology in the form of an abstraction hierarchy using Formal Concept Analysis (FCA) [45, 46]. MCRDR incrementally and rapidly acquires and validates the knowledge on a case-by-case basis. Cases are used to motivate and validate the knowledge acquired. Figure 1 shows the process involved. MCRDR use cases found naturally in a domain to assist the user with classification of those cases and development of rules to cover the cases. The technique is very simple. Classification of cases and rule development are part of the same task. An expert runs an inference on a case using the current knowledge in the knowledge base (KB). If the expert agrees with the system assigned conclusion they review another case. If they disagree they assign the correct conclusion, which may involve adding that conclusion to the set of valid conclusions. Next the expert picks some features in the case, which provide a justification of the new conclusion and form the rule conditions. The rules are then automatically translated into crosstable format. FCA uses the crosstable to generate higher-level concepts by taking intersections of shared rule conditions. The concepts are ordered

using subsumption (≤) to generate a complete lattice. It is also feasible to use records in a database as the basis of the crosstable. Ideally a machine learning algorithm should be used on the dataset prior to translation to a crosstable to remove irrelevant attributes. This is not necessary when rules are used as input since the relevant features have been identified by the expert during knowledge acquisition. The ontology developed is based on validated knowledge, real examples and there is strong ontological commitment between the performance and explanation system. Let us examine reverse engineering, MCRDR and FCA in more detail.

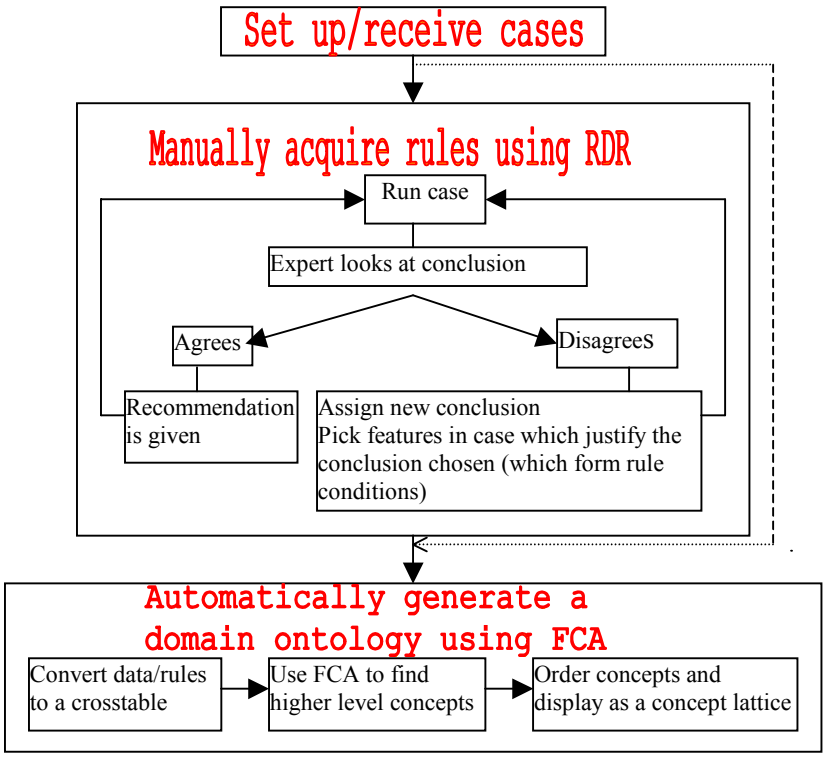

Figure 1: A process model for reverse engineering a domain ontology. MCRDR will classify cases and identify salient features before application of FCA. Alternatively, examples/cases can be used directly as input into the final phase. As KA in MCRDR is incremental and on a case-by-case basis the process is iterative.

2. Reverse Engineering an Ontology

Reverse engineering has been described as uncovering 'secrets' in the original artifact not apparent in the development of the product [26]. Unlike the majority of knowledge-based systems (KBS) approaches, the development of an MCRDR KBS does not require the user, typically the domain expert, to structure their knowledge, specify relationships between concepts or to provide abstract concepts. We can therefore view the retrospective discovery of an ontology from an MCRDR KB as reverse ontological engineering. As described above, in the MCRDR approach the

domain expert enters rules in response to reviewing system recommendations for cases. The conclusion and rule conditions (features in the case) become an exception rule to the rule which gave the misclassification. To assist the user and provide online validation, cases which prompt a rule to be added are stored in association with the new rule (these are known as cornerstone cases) and shown to the user when an exception rule needs to be added. The features selected to form the exception rule must distinguish between the current case and the cornerstone case(s) for the rule which misfired. See Figure 2 for an example of a partial MCRDR KB. The expert may also stop an existing rule and add a new rule at the first level if they do not agree with the features in the parent rule. In standard rule-based systems, validation of the entire knowledge base may be needed each time one rule is added [36]. From experience with over a dozen currently deployed systems and through experimentation we have found that the time to acquire and validate a rule in an MCRDR KB is constant at approximately one rule per minute. Most of that time is taken up with deciding on the preferred wording of the conclusion. The order of cases seen will affect the compactness and time to maturity of an RDR KB but does not noticeably impact inferencing time. When FCA is applied the KB is restructured using term subsumption and the order is even less of an issue. MCRDR have found commercial success in the domains of pathology report interpretation (LabWizard [19]) and help-desk applications [15]. Knowledge bases with over 7,000 rules have been developed in less than a month. Key features of MCRDR that have assisted easy knowledge acquisition are the use of:

1. an exception structure for knowledge representation which provides local patching of the rules; and
2. the use of cases to motivate, assist and validate the acquisition of new knowledge.

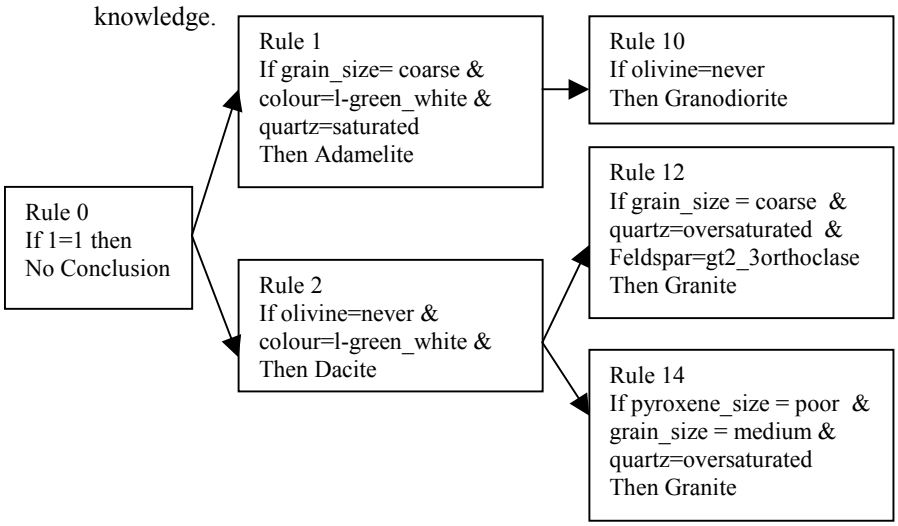

Figure 2: A partial MCRDR for the domain of Igneous Rocks. Two levels of decision lists are shown. The children of true parents are evaluated against the case. The last true rule on each pathway is reported as the conclusion/s.

Once we have the performance system we use FCA to uncover the structure and higher-level concepts. A concept in FCA is comprised of a set of objects and the set of attributes which are shared by those objects. In our usage of FCA we treat each rule as an object (identified by its rule number and conclusion code) and the rule conditions as the attributes [31]. In Figure 3 we have translated three of the rules in Figure 2 into a crosstable by picking up the rule conditions from all ancestors. We have selected the rules for the granite family of rocks, but we could have developed an ontology including all 6 rules in Figure 2 (or the whole KB). This representation is known as a formal context in FCA and is used to generate concepts. FCA works by finding intersections of rule conditions in different rule objects. These intersections represent higher-level concepts shared by the primitive concepts. The ability to show higher levels of abstraction than were originally specified means that the lattice does more than just restructure or re-represent the original low-level concepts. The terms, structure and coverage of the ontology will depend on the content of the rules (which are directly related to the content of the cases), or the content of the data records if mining from databases. Once the concepts have been generated they are ordered using term subsumption to produce a concept lattice. The concept lattice is visualized as a line diagram. The line diagram based on the formal context in Figure 3 is given in Figure 4. From Figure 4 we can identify what features these three types of granite share and what features distinguish them from one another. To find attributes that belong to a concept follow all ascending paths. Objects sharing an attribute are reached by descending paths. In the supremum (top concept) we can see that they share the features COLOUR is LIGHT_GREEN_WHITE and OLIVINE is NEVER. They are further distinguished from one another by other attributes. We can see that granodiorite %GR000 shares the attribute GRAIN_SIZE=COARSE with granite %GR001 but that they differ in their QUARTZ content, with granodiorite having QUARTZ=SATURATED and granite having QUARTZ=OVER-SATURATED. Granite is similar to microgranite %GR002 in the QUARTZ content but differs in GRAIN_SIZE which is MEDIUM. Granite also takes into account the FELDSPAR content and tbe rule for microgranite includes PYROXENE=POOR. The diagram raises the question of whether there a relationship between PYROXENE and FELDSPAR. If there is rules 12 and 14 may be considering the same features.

	1=1	Grain-size= coarse	Colour=light _green_white	Colour=light _green_white	Quartz =saturated	Olivine= Never	Quartz =over- saturated	Feldspar= GT2-3 Orthoclase	Grain-size= medium	Pyroxene= poor
10-%GR000	X	X	X	X	X	X				
12-%GR001	X	X	X	X		X	X	X		
14-%GR002	X		X	X		X	X		X	X

Figure 3: A formal context for the rules 10-%GR000(Grandiorite), 12-%GR001(Granite) and 2-%GR002(Microgranite)

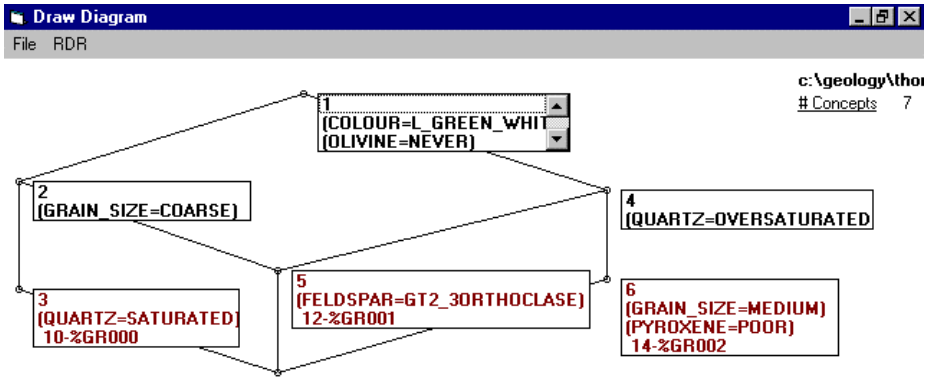

Figure 4: The line diagram for the formal context in Figure 3 showing the relationship and structure between rules 10-%GR000(Grandiorite), 12-%GR001(Granite) and 14-%GR002(Microgranite).

Our ability to generate a concept lattice from rules is not restricted to tasks of classification or MCRDR KBs. Most RDR have been developed for classification tasks but they have also been applied to causal modeling, configuration, design, simulation and search control problems. The same knowledge acquisition and representation technique has been applied to these problems but the inference engine has been modified in some cases, generally requiring additional cycles and sometimes input from the user for conflict resolution between alternatives. However, the generation of an ontology from the rules is the same regardless of the type of domain or task. Work has also been done which looks at the generality of using FCA for creating concept lattices from any propositional knowledge representation that can be mapped into a crosstable, known as a formal context. The line diagram shown in Figure 5 uses rules from an animal knowledge base shipped with CLIPS 5.1. The use of FCA to generate lattices from databases is well established and used in a number of large companies [43].

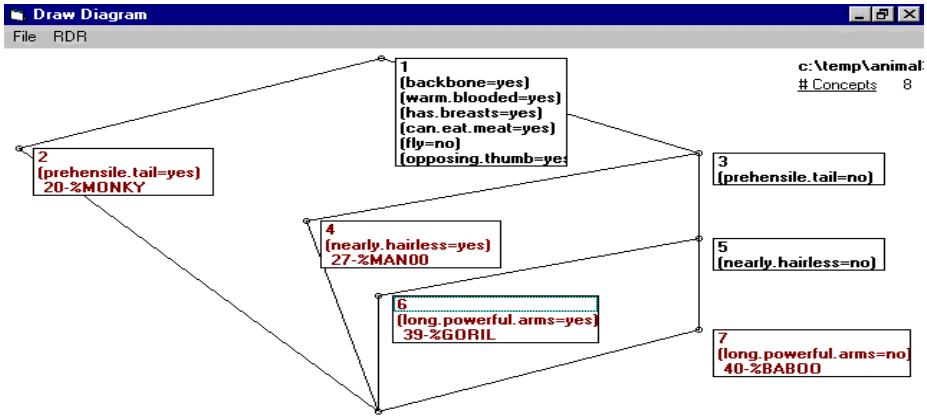

Figure 5: Finding the animals closest to man using CLIPS 5.1 Animal KBS. From the line diagram we can see that the monkey is separated the most from man and the gorilla and baboon are the most similar.

3. To Engineer or Reverse Engineer

The goals of retrospective discovery of an ontology are similar to the goals of ontological engineering and reverse engineering. Ontological engineering supports reuse and sharing of knowledge. Although we do not need the ontology for knowledge acquisition using MCRDR, we are interested in reusing the knowledge acquired for inferencing or consultation for alternative uses such as tutoring, modeling, explanation, critiquing and 'what-if' analysis. Since we go straight from data to rules without the development of an intermediate model we need to reverse engineer models.

The notion of cost-saving using reverse engineering of knowledge bases may seem absurd if we view ontologies as a solution to the problem of acquiring knowledge. Ontologies are seen as a means of obtaining a wealth of structured, and hopefully already validated, knowledge that can be used in the development of KBs. Unfortunately, many of the difficulties associated with the KA bottleneck, such as the unwillingness or the inability of experts to articulate their knowledge, difficulty in understanding and formalizing the conceptual model of another person, the lack of agreement on terms and the unreliability and variability of models, also hamper ontology acquisition. We seem to have a chicken and egg problem. To address the ontology acquisition bottleneck, there is a strong focus on the sharing and reuse of ontologies. One approach is the development of libraries of reusable ontologies containing components which could be assembled together. However, integration of ontologies is itself seen "as a challenging task" [22, p.53]. This problem occurs due to the content, partitioning and structure of concepts and the variability in the knowledge representation used. An even bigger impediment to the reusable library of ontologies approach appears to be the *interaction problem,* which contests the notion that problem solving knowledge and domain knowledge are two distinct components of a KBS. These two types of knowledge are difficult to separate since the domain knowledge needed will be strongly affected by the type of problem and the method of inference [3]. None of these approaches are an easy or complete solution to the high cost of developing ontologies. Russ et. al. [33] offer a discussion of the tradeoffs between reusability and usability of ontologies and conclude that current tools do not provide a practical solution.

Noy and Musen [23] offer two solutions to the problem of integration:

1. merging multiple ontologies into one;
2. aligning ontologies in a way that allows sharing and reuse of information between them.

The latter can be assisted by adopting specific principles for the development of a *core ontology* [41]. Like any standard gaining acceptance amongst competing alternatives is difficult.

The ability to generate a term-subsumption hierarchy using our approach could possibly be used as a framework for comparing and combining existing ontologies, or parts thereof, similar to the approach proposed by [44]. We have developed a framework based on MCRDR and FCA which allows multiple knowledge bases to be combined [30]. Describing a concept using an intensional and extensional definition assists in determining whether terms are being used in a similar way. For example, Figure 6 shows a line diagram which includes the rules for classifying the

rock Diorite (coded as %DI000) from 5 different KB, identified as C1, C2, C3, C5 and L3. While we don't have space to explore all that can be learnt from this lattice, we draw attention to the different ways in which the feature grainsize is used. We can see that experts C3 and C1 consider the grainsize of Diorite to be coarse. C2 and L3 consider the grainsize to be medium. This difference in perception is made explicit in the lattice. A difference in terminology is shown in concept 6, where expert C5 considers the grainsize for Diorite to be coarsely-crystalline. This appears to match the view of experts C1 and C3. Another difference in terminology could be [Olivine=No] (concept 7 expert L3) and [Olivine=Never] (concept 8 expert C1). We reconcile similar terms by using a table for mapping synonyms, hypernyms and hyponyms. The use of mappings is a common (e.g. [9]) alternative approach to the use of core ontologies suggested above [41]. The framework we have developed includes conflict detection, negotiation and resolution strategies. We determine which resolution strategy to use based on a four-state model of comparison [34]. See [30] for details regarding the reconciliation framework.

Another feature shown in Figure 6 is the ability to provide a case in support of a rule. This is particularly useful for conflict resolution. In the KA technique offered by standard FCA (that is FCA without MCRDR) 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 6, the case associated with concept number 8 has been popped up. This rule states that Expert C1 believes IF [Colour=Mesocratic], [Olivine=Never], [Quartz=Possibly], [Grain_Size=Coarse] THEN the rock is Diorite (%DI000). 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 C1 expert from having to think up, find or recall an example.

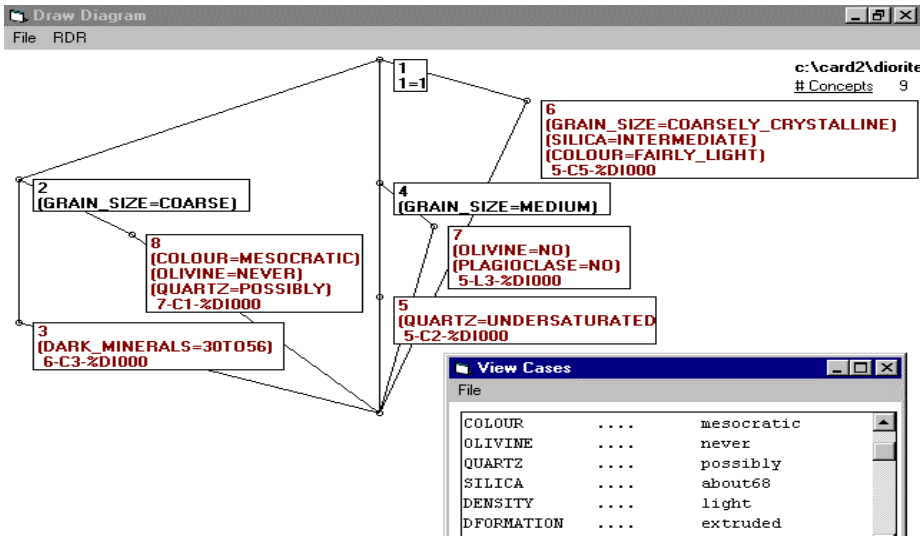

Figure 6: A line Diagram showing how different KB can be compared and how a case can be popped up to assist in reconciliation of conflict between multiple sources of expertise.

Apart from savings in time and effort, automatic generation of an ontology from a rule-base reduces validation overheads as there is high fidelity between the two. We avoid the interaction problem by concentrating on contextualised domain knowledge. The ontological commitment made is directly related to the cases that have been seen. Also, the rules represent performance knowledge based on the visible and everyday behaviour of the domain expert (which is: see case, assign conclusion, select features to justify conclusion) and may be executed for validation with real cases. Rather than starting with models, which by their very nature are imperfect representations, we start with real cases for acquisition and validation of the rules.

4. A Comparison to other Approaches

Most approaches are founded on the basic assumptions that the cost of knowledge acquisition is higher than the cost of building the ontology and that KA will more easily follow once the ontology is available. While progress is being made on a number of fronts, the MCRDR/FCA approach challenges these assumptions and allows ontology building to follow knowledge acquisition in domains where cases and a domain expert exist.

The current use of ontologies in the MCRDR/FCA approach is for the purpose of modelling domain knowledge. GenSim [16] shares this purpose however their main focus is on simulation and prediction of experimental results and GenSim is restricted to the domains of molecular biology and biochemistry. As previously mentioned, the motivation for adding FCA techniques into MCRDR was the desire to provide an environment that allowed knowledge to be reused in multiple ways for various activities and decision styles. Other similarly motivated work are Protégé [28] and the work on reusing the knowledge in MYCIN [4], however both approaches targeted a particular reuse activity such as automatic development of screens for KA or tutoring, respectively.

A key difference between the ontology developed using FCA and a semantically organised ontology such as WordNet [21] is the use of term subsumption for concept structuring. While WordNet uses a hierarchy of superclasses and subclasses this hierarchy, known as synsets, is based upon the sense or meaning of the word. However, the FCA ontologies are not simply syntactical representations. FCA is focused at the concept level and provides the connection between the names and definitions at the representation level and the objects at the object level. FCA restructures the RDR knowledge based on the terms used, not based on the order of rules. When we describe a conclusion by its rule conditions we are describing the meaning of that conclusion. Most ontologies focus on the intensional description of a concept, but the philosophically based notion of a formal concept defined by its intension and extension in FCA adds meaning and assists in comparison of concepts.

Another distinguishing feature between FCA and WordNet, or other lexically-based ontologies, is that FCA does not develop one or more hierarchies which are broken up into nouns, verbs, adjectives, etc. The type of terms which appear in a concept lattice are directly dependent on the type of terms used in the formal context. The Generalized Upper Model (GUM) [1] provides a level of abstraction which is in between lexical knowledge and conceptual knowledge. It may be that some benefit

can be achieved by including a lexically-based ontology into the FCA ontology so that semantically equivalent terms can be identified and possibly automatically reconciled. Unlike GUM, which has separate hierarchies for concepts and relationships, the FCA concept lattice combines concepts and relationships into the one structure. The combination seems to be more visually comprehensible.

The use of taxonomies in the organization of ontologies is common [22]. The structure of these taxonomies varies with some approaches having one large taxonomy e.g. CYC [12], Sowa's ontology [37], GUM, GenSim, and noun sysnets in WORDNET, or a number of smaller ontologies e.g. TOVE [11]. In some sense MCRDR/FCA offers both approaches. The user may choose to develop a concept lattice of the whole KB or a collection of KBs or many smaller lattices can be generated based on restricted contexts (or rules subsets). Due to the limitations of displaying and navigating around large graphs and computational complexity issues the latter approach is recommended and we thus offer various techniques for selecting rules/conclusions to include in a formal context. The FCA concept lattice is closest to the implicit taxonomy found in PLINIUS that is structured using the subsumption relation. The PLINIUS Project [42] does not use the hierarchical-axiomatic approach used by most other projects for building ontologies. At a surface-level the construction of the PLINIUS ontology is similar to the FCA technique in that it begins with a set of atomic concepts. Elements of these sets are combined to define other concepts. PLINIUS uses rules in its construction kit for determining sub and superconcepts rather than term subsumption. PLINIUS is very domain specific and it is not known if it can be extended beyond the domain of chemistry. Although the FCA concept lattices tend to be domain specific this feature is due to the nature of the knowledge contained in the KBS. There is no particular restriction on the domain that they can describe. Two further significant differences are that the FCA concept lattice is automatically generated from the primitive concepts and is able to represent properties belonging to an object.

In addition to containing a taxonomy of concepts, most ontologies include a set of meaningful properties and categories [22]. Some ontologies will allow explicit specification of axioms. This is not possible in an MCRDR KB. The absence of axioms in MCRDR and FCA reflects the strongly held view of knowledge applying within the context of a case or example. MCRDR is able to handle knowledge which goes beyond first-order logic but this knowledge is not always clearly displayed in the concept lattice. For example a default rule will appear in the top concept in the lattice but it is not labeled or otherwise identified as a default. Similarly, context is handled in MCRDR through its exception structure and the storage of cases which prompt new knowledge to be added. As shown in Figure 6, cases associated with a particular node on the lattice may be shown and the exception structure can be found in the lattice although it may not be as obvious as in the original MCRDR format (compare Figures 2 and 4).

There are a number of ontological approaches such as [8, 40] that concurrently perform KA and develop the ontology. The Methontology [8] ontology-development process includes specification, acquisition and conceptualization of domain knowledge, integration with existing ontologies, evaluation and automatic implementation into Ontolingua. In such approaches, ontological engineering is not a prerequisite task to KA but a complementary task that can be used to assist and

validate the knowledge being acquired. What makes MCRDR/FCA so unusual is that KA does not involve the specification of the user's conceptual model and the ontology is truly automatic and can only be generated after KA has occurred. Similar to the approach offered by ROCHS, mentioned below, it may be desirable to integrate the MCRDR/FCA approach with another ontological approach to improve the systematic acquisition of knowledge or to speed up the growth of the KB. However, since MCRDR is fundamentally an incremental KA technique designed to deal with cases as they occur (and since the current approach results in rapid development anyway) such changes to the approach may be counterproductive.

Ontologies are used in data mining by a number of approaches. In some approaches the output of data mining is an ontology. Omelayenko [24] has surveyed a number of approaches to learning ontologies. He divides the approaches broadly into two main categories depending on whether they involve machine-learning (ML) or manual construction. Omelayenko found, and we agree, that the combination of both techniques tended to offer the best results by combining the speed of machine learning with the accuracy of a human. Omelayanko also found that ML results in flat homogeneous structures often in propositional form. Flat structures do not support queries at multiple level. We have proposed [32] using FCA on rule bases (possibly output by other datamining techniques) to develop these multilevel rules.

Others have used ontologies to guide the discovery of multi-level rules (e.g the ParkaDB approach [39]. In ParkaDB an ontology (in the form of a concept hierarchy) together with frequency counts are used to determine which concepts should be included in a rule. The ontology provides the background knowledge to guide the discovery process. A number of similar approaches that use concept taxonomies (e.g. [35]) have also been developed but these are based on traditional relational database technology and require transformation to a generalized table as part of the preprocessing which can result in over-generalisation. DeGraaf et al. [10] use a prespecified taxonomy to suggest interesting association rules, which they define as unexpected rules that do not fit with the taxonomy. Intuitively it is appealing but the approach obviously relies on the development of a complete and valid taxonomy. ParkDB does not require such preprocessing but supports dynamical generalization of data without over generalization. Our approach is different to all of these in that we do not use a concept hierarchy to develop rules. Instead we use rules to develop a concept hierarchy which may lead to higher level rules being uncovered. Thus we avoid the substantial effort required in first developing the hierarchies and the difficult task of validating them. Given any string or substring and using set intersection, term subsumption and lectical ordering we are able to find all combinations using that string. Some concepts will include attributes and/or objects at different levels of abstraction such as the objects mammal and dog will appear in individual and shared concepts. The combination of multi-level concepts allows us to perform queries at and across multiple levels.

The connection between FCA and ontologies was first made in the work in Web Analysis and Visualisation Environment (WAVE) [17] which defines an "ontology as a specification of a concept lattice". However, our application of FCA to rules bases is novel. Other MCRDR work using conceptual hierarchies includes: Ripple-down rules Oriented Concept Hierarchy System (ROCHS) [20], Nested MCRDR (NRDR) [2] and the discovery of class relationships from MCRDR KBS [38].

5. Conclusion and Future Work

RDR and MCRDR have been evaluated empirically [7] and experimentally [6]. Of greater concern to the approach described in this paper is the evaluation of the role and use of ontologies as part of the MCRDR approach. Some evaluation has been done concerning the value of the concept lattice for explanation, exploring and learning about a domain [29]. Part of this evaluation involved obtaining four different MCRDR KB from the fields of pathology, agriculture, geology and chemistry. The MCRDR rules were used to develop FCA lattices. The lattices were used by a beginner (level lower than novice) to answer questions and learn about the domain. The knowledge 'learnt' after about 1 hour of generating and browsing the lattices was written down and shown to domain experts who were asked to comment on the validity and value of what had been learnt. In the first 3 domains each expert found that the understanding gained was valid for that domain. Experts from the pathology and agricultural domains were impressed by the depth of understanding which the beginner appeared to have. In the geology and pathology case studies, where the beginner and expert were able to interact, the lattices provided a valuable communication channel for discussing key ideas and modifying hypotheses. The rules in the chemistry domain tended to have single conditions which resulted in few intersections between rules and uninteresting lattices so there was very little in terms of higher level concepts, structure or relationships to learn from that knowledge source. While such an experiment is not conclusive it does indicate that the ontologies represented in the concept lattice offers knowledge at a greater depth which extends beyond offering an alternative graphical view of the rules.

The work by Predigger [25] which takes a many-valued context and derives concept graphs which show on the edges the nature of the relationships between the vertices would provide an even richer representation. Just as it is necessary in the MCRDR/FCA approach for a human to initially provide the label for any abstractions uncovered, the nature of the relationship, apart from the subsumes relation, would need to be supplied as this information is not automatically derivable from the KB. The work by Priss and Old [27] which extends FCA to Relational Concept Analysis (RCA) considers relationships other than subsumption will be investigated as a possible means of generating and displaying more kinds of relationships between concepts. As suggested in the Section 4, if the MCRDR/FCA approach were extended to allow the use of other ontologies such as WordNet it may be possible to automatically label the edges with the relationship type. Integration is not an easy task. Asking the user to supply relationship names could be more efficient and accurate. Another lesser problem is that the labeling of edges will decrease the number of concepts that can be shown before comprehensibility is lost. Increased expressibility in a representation always comes with a cost.

It would be interesting to explore the ramifications of ontologies for MCRDR from a theoretical standpoint. Sowa's ontology is philosophically motivated and seeks to address a number of theoretical issues concerning ontologies. The similarity between Sowa's fundamental principles of ontology design as distinctions, combinations, and constraints [37] is embodied in the FCA set-theoretic approach and the use of difference lists, cornerstone case lists (which identify related cases)

and rules in MCRDR. A key difference between the constraints used by Sowa and those found in MCRDR KB is that in Sowa's ontology the ontology developer is responsible for identifying logical constraints. In MCRDR, the use of cases constrains what can be represented in the rules.

As mentioned before, building an ontology from an RDR KB means that we are starting with validated knowledge. Validation of each rule extends beyond the cornerstone case. Each time a case is seen it is validated against the whole KB. If any case is deemed to be incorrect based on that rule, an exception rule is added. In our commercial version of MCRDR we record how many cases have visited a rule and been accepted. Selection of which rules to include in the generation of an ontology could be made on the basis of these statistics. A rule that has been visited by many cases can be seen to be more important and more reliable than one that has rarely been exercised or that has been frequently patched.

There are many open issues in the field of ontological engineering. Until definitive answers are found, alternatives such as that offered by MCRDR/FCA deserve some consideration. The alternative offered is to rapidly acquire knowledge and use that knowledge to automatically reverse engineer an ontology.

Acknowledgments
RDR research is supported by Australian Research Council grants.

References

[1] Bateman, J.A., Magnini, B. and Rinaldi, F. (1994) The Generalised Italian, German, English Upper Model. *Workshop on Comparison of Implemented Ontologies at 11th European Conf. on Artificial Intelligence (ECAI'94),* Aug.8-12, Amsterdam, The Netherlands.

[2] Beydoun, G. and Hoffmann, A. (1997) Acquisition of Search Knowledge, In Enric Plaza and Richard Benjamins (eds), *Knowledge Acquisition, Modeling and Management. 10th European Workshop, EKAW'97*, LNAI 1319, Springer-Verlag, 1-16.

[3] Bylander, T. and Chandrasekaran, B. (1988) Generic Tasks in Knowledge-Based Reasoning: The Right Level of Abstraction for Knowledge Acquisition In B.R. Gaines and J.H. Boose (eds.) *Knowledge Acquisition for KBSs*, Academic Press, London.

[4] Clancey, W.J., (1986) From GUIDON to NEOMYCIN and HERACLES in Twenty Short Lessons: ORN Final Report 1979-1985 *AI Magazine*, August 1986, 40-60.

[5] Compton, P. and Jansen, R., (1990) A Philosophical Basis for Knowledge Acquisition. *Knowledge Acquisition* 2:2

[6] Compton, P., Preston, P. and Kang, B. (1995) The Use of Simulated Experts in Evaluating Knowledge Acquisition, *Proceedings 9th Banff Knowledge Acquisition for Knowledge Based Systems Workshop* Banff. Feb 26 - March 3 1995.

[7] 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.

[8] Fernandez Lopez, M., Gomez-Perez, A., Pazos Sierra, A., and Pazos Sierra, J. (1999) Building a Chemical Ontology Using Methontology and the Ontology Design Environment *IEEE Intelligent Systems* 14(1) January/February 1999.

[9] Gennari, J.H., Tu, S.W., Rothenfluh, T.E. and Musen, M.A. (1994) Mapping Domains to Methods in Support of Reuse Int. Journal of Human-Computer Studies 41:399-424.

[10] Graaf, de, J., M., Kosters, W., A. and Witteman, J., J., W. (2000) Interesting Association Rules in Multiple Taxonomies, Proc. of the 12th Belgium-Netherlands Artificial Intelligence Conf.

[11] Gruninger, M. and Fox, M.S. (1995) Methodology for the Design and Evaluation of Ontologies *Workshop on Basic Ontological Issues in Knowledge Sharing* 19-20 Aug., Montreal, Quebec.
[12] Guha, T.V., and Lenat, D.B. (1990) CYC:A Mid-Term Report *AI Magazine* 11(3):32-59.
[13] Jasper, R. and Uschold, M. (1999) A Framework for Understanding and Classifying Ontology Applications In *Proc. of the 12th Workshop on Knowledge Acquisition, Modeling and Management (KAW'99),* 16-21 October, 1999, Banff.
[14] 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.
[15] Kang, B., Yoshida, K, Motoda, H, and Compton, P. (1997) Help Desk with Intelligent Interface *Applied Artificial Intelligence* 11:611-631.
[16] Karp, P.D. (1993) A Qualitative Biochemistry and its Application to the Regulation of the Trytophan Operon In L. Hunter (ed.) *Artificial Intelligence and Molecular Biology* Menlo Park, AAAI Press, California: 289-325.
[17] Kent, R. E. and Neuss, C. (1995) Creating a 3D Web Analysis and Visualization Environment *Computer Networks and ISDN Systems*, 28:109-117.
[18] Kietz, J.-U., Maedche, A. and Volz, R. (2000) A Method for Semi-Automatic Ontology Acquisition from a Corporate Intranet. WS "Ontologies and Text", co-located with *EKAW'2000,* Juan-les-Pins, French Riviera, October 2-6, 2000.
[19] Lazarus, L. (2000) Clinical Decision Support Systems: Background and Role in Clinical Support, http://www.pks.com.au/CDSS_White_Paper_doc.pdf
[20] Martinez-Bejar, R., Benjamins. R., Compton,, P., Preston, P. and Martin-Rubio, F. (1998) A Formal Framework to Build Domain Knowledge Ontologies for Ripple-Down Rules-Based Systems. In Gaines, B. and Musen, M. (eds) *Eleventh Workshop on Knowledge Acquisition, Modeling and Management (KAW'98),* Banff, Alberta, Canada, 18-23 April, 1998, Vol 2:SHARE.13.
[21] Miller, G. A. (1990) WORDNET: An On-Line Lexical Database, *Int. Journal of Lexicography* 3-4:235-312.
[22] Noy, N. F. and Hafner, C. D. (1997) The State of the Art in Ontology Design: A Survey and Comparative Review *AI Magazine* Fall 1997 18(3): 53 - 74:
[23] Noy, N. F. and Musen, M. A. (1999) SMART: Automated Support for Ontology Merging and Aligning In *Proceedings of the 12th Workshop on Knowledge Acquisition, Modeling and Management (KAW'99),* 16-21 October, 1999, Banff.
[24] Omelayenko B., (2001) Learning of Ontologies for the Web: the Analysis of Existent Approaches, In: Proceedings of the International Workshop on Web Dynamics, held in conj. with the 8th Int. Conf. on Database Theory (ICDT'01), London, UK, 3 Jan. 2001.
[25] Predigger, S. (1999) The Lattice of Concept Graphs of a Relationally Scaled Context In W. Tepfenhart and W. Cyre *Proceedings fo the 7th International Conference on Conceptual Structures; Standards and Practices (ICCS-99)* Blacksburg, VA, USA, July 12-15, LNAI-1640, Springer Verlag, Berlin: 401-414.
[26] Pressman, R. (2000) 5th Edn. *Software Engineering: A Practitioner's Approach* (European Adaptation), McGraw Hill, Berkshire.
[27] Priss, U. and Old, J. (1998) Information Access through Conceptual Structures and GIS In *Information Access in the Global Information Economy Proceedings of the 61st Annual Meeting of ASIS'*98 :91-99 also http://php.indiana.edu./~upriss/papers/asis98.ps
[28] Puerta, A. R, Egar, J.W., Tu, S.W. and Musen, M.A. (1992) A Multiple Method Knowledge Acquisition Shell for Automatic Generation of Knowledge Acquisition Tools *Knowledge Acquisition* 4(2).
[29] Richards, D. (1998) An Evaluation of the Formal Concept Analysis Line Diagram In Slaney, J., Antoniou. G and Maher, M.J. (eds) *Poster Proc. of 11th Australian Joint Artificial Intelligence Conf. AI'98* 13-17 July 1998, Brisbane, Australia, 121-133.

[30] Richards, D. (2000) Reconciling Conflicting Sources of Expertise: A Framework and Illustration In P. Compton, A. Hoffmann, H. Motoda and T. Yamaguchi1 (eds) *Proceedings of the 6th Pacific Knowledge Acquisition Workshop* Sydney December 11-13,2000, 275 – 296.

[31] 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), *Conceptual Structures: Fulfilling Peirce's Dream, Proceedings of the Fifth International Conference on Conceptual Structures (ICCS'97),* August 3 - 8, University of Washington, Seattle, USA, LNAI 1257, Springer Verlag, Berlin, 198-212.

[32] Richards, D. and Malik, U. Multi-Level Knowledge Discovery in Rule Bases, *Applied Artificial Intelligence*, Taylor and Francis Ltd. (accepted 2002).

[33] Russ, T., Valente, A., MacGregor, R. and Swartout, W. (1999) Practical Experiences in Trading Off Ontology Usability and Reusability In *Proceedings of the 12th Workshop on Knowledge Acquisition, Modeling and Management (KAW'99),* 16-21 Oct., 1999, Banff.

[34] Shaw, M.L.G. and Gaines, B.R., (1988) A Methodology for Recognising Consensus, Correspondence, Conflict and Contrast in a Knowledge Acquisition System in *Proceedings of the 3rd Knowledge Acquisition for Knowledge-Based Systems Workshop,* November 1988, Banff, Canada.

[35] Singh, L., Scheuermann, P. and Chen, B. (1997) Generating Association Rules from SemiStructured Documents Using an Extended Concept Hierarchy. In *Proceedings of the 6th International Conference on Information and Knowledge Management*. Las Vegas, Nevada, USA, 1997, pp. 193-200.

[36] Soloway, E, Bachant, J. and Jensen, K. (1987) Assessing the Maintainability of XCON-in-RIME: Coping with Problems of a very Large Rule Base *Proceedings of the Sixth Int. Conf. on Artificial Intelligence*, Seattle, WA: Morgan Kaufman, Vol 2:824-829.

[37] Sowa, J. F. (1995) Top-Level Ontological Categories *International Journal of Human-Computer Studies* 43(5-6): 669-685.

[38] Suryanto, H. and Compton, P. (2000) Discovery of Class Relations in Exception Structured Knowledge Bases In *Proceedings of the 8th International Conference on Conceptual Structures; Logical, Linguistic and Computational Issues (ICCS-2000)* Darmstadt, Germany, 14-18 August, Springer Verlag, Berlin.

[39] Taylor, M., Stoffel, K. and Hendler J. (1997) Ontology-based Induction of High Level Classification Rules In *SIGMOD Data Mining and Knowledge Discovery workshop proceedings*. Tuscon, Arizona, 1997.

[40] Uschold, M. and Gruninger, M. (1996) Ontologies: Principles, Methods and Applications. *Knowledge Engineering Review* 11(2):93-155.

[41] Valente, A. and Breuker, J. (1996) Towards Principled Core Ontologies In B.Gaines and M. Musen (eds.) *Proceedings of the 10th Workshop on Knowledge Acquisition, Modeling and Management (KAW'99),* Banff.

[42] Van de Vet, P. E., Speel, P. -H. and Mars, N.J. (1994) The PLINIUS Ontology of Ceramic Materials. 11th European Conf. on Artificial Intelligence (ECAI'94) *Workshop on Comparison of Implemented Ontologies*, 8-12 Aug., Amsterdam, The Netherlands.

[43] Vogt, Frank and Wille, Rudolf (1995) TOSCANA - A Graphical Tool for Analyzing and Exploring Data *Knowledge Organisation* 22(2):78-81.

[44] Weinstein, P. and Birmingham, W. P. (1999) Comparing Concepts in Differentiated Ontologies In *Proceedings of the 12th Workshop on Knowledge Acquisition, Modeling and Management (KAW'99),* 16-21 October, 1999, Banff.

[45] Wille, R. (1982) Restructuring Lattice Theory: An Approach Based on Hierarchies of Concepts In Ordered Sets (Ed. Rival) pp:445-470, Reidel, Dordrecht, Boston.

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