The Knowledge Engineering Review, Vol. 6:2, 1991, 97–120

Case research on knowledge acquisition: observations and lessons

CHRISTINE CHAN and IZAK BENBASAT*

School of Computing Science, Simon Fraser University, Burnaby, BC, Canada V5A 1S6

Abstract

Expert systems are being built despite the widely acknowledged problem of acquiring knowledge from experts. This study attempts to understand how knowledge acquisition is conducted in practice by investigating three expert system development projects. A CASE research methodology is adopted, and data is collected through unobtrusive observation, from taped protocols of knowledge acquisition sessions, retrospective interviews with the participants involved, and deliverables produced. The variables examined include the problem domain, the domain expert, the knowledge engineer, the knowledge acquisition process, the expert system construction process, potential users, organizational setting, and the expert system itself. The knowledge acquisition processes for three expert systems in the domains of law of negligence, telephone line fault diagnosis, and wastewater treatment have been examined. By juxtaposing the observations drawn with findings from the relevant literature, the study makes prescriptive suggestions on considerations and techniques for future acquisition efforts, and provides data for hypothesis generation in further research.

1 Introduction

Knowledge Acquisition (KA) is the general name given to the process of eliciting, acquiring, and representing knowledge consisting of descriptions, relationships and procedures in a specialized domain of interest. KA is considered to be a major bottleneck in the efficient development of expert systems (see, for example, Breuker *et al.*, 1987*a*; Abrett *et al.*, 1987). While the difficulty of accessing expertise is widely acknowledged in the literature, expert systems are being built at an increasing rate. Dupont alone projects that by the year 1991, over 2000 systems will be in place (Watkins *et al.*, 1987). It has been pointed out that most knowledge engineers use an "intuitive" approach to knowledge acquisition (Alexander *et al.*, 1987). This intuitive approach has, in many cases, succeeded in sufficiently covering the problem domain and providing a basis for expert systems that perform satisfactorily. At the same time, substantial efforts to developing a set of KA techniques and automated tools are taking place in order to elevate KA to a "science" rather than an art (Welbank, 1987; Boose, 1989). However, most automated tools are currently at a prototype stage, and the quality of their interfaces is not yet adequate (Kitto, 1988).

The objective of this study is to investigate the methods and techniques of KA used by practitioners in the field. The need for empirical research into expert systems development has been mentioned by several authors. Gaines advocated more variety, detail and evaluation in "the case histories of applications of the techniques and tools" (Gaines, 1987b). Referring to the lack of understanding of expert system implementation issues, Sviokla states that :

> "... more empirical evidence is needed from which more generalizable inferences can be drawn. Experience is still scant and generalizations are hard to create with confidence." (Sviokla, 1986, p 363).

*Faculty of Commerce and Business Administration, University of British Columbia, Vancouver, Canada V6T 1Y8

This study attempts to address this need for empirical studies by investigating how the process of knowledge acquisition is implemented in expert system development. It is expected that such an inquiry would shed light upon the process by providing a set of useful observations to guide future acquisition efforts on a practical level and data for hypothesis generation for research purposes.

The paper proceeds as follows: section 2 describes some relevant literature, and puts the present inquiry into the context of some existing writings on KA; section 3 explains the research strategy and design; section 4 presents the data collected from the three cases; section 5 discusses some observations drawn from the cases, their comparisons with notions about KA suggested in the literature as well as some limitations of the study; and section 6 presents some concluding remarks.

2 Background Literature

At least two sets of literature explicitly address the topic of KA. The first explains and evaluates the techniques and tools associated with KA; the second does not deal with the detailed techniques involved, but describes the KA process as a phase within the expert system development life cycle.

KA has been defined in Anjewierden (1987) and Breuker and Wielinga (1987*b*). The process is subdivided into the three stages of knowledge elicitation, knowledge representation, and knowledge encoding (Diaper, 1989), or the areas of knowledge collection, knowledge analysis, and knowledge representation (Bell and Hardiman, 1989). Most of the techniques and methods of KA in the first set of literature address the first stage of knowledge elicitation or collection. In-depth reviews and evaluations of the techniques, tools and methodologies are found in Boose (1989), Burton and Shadbolt (1987), Cordingley (1989), Hart (1986), Olson and Reuter (1987), Neale (1988), Greenwell (1988) and Wellbank (1987). A categorization scheme to classify these tools and techniques has been suggested by Dhaliwal and Benbasat (1990):

1 Automated *versus* non-automated tools, i.e., the degree to which the technique is implemented as a computer program.
2 Direct methods which require experts to report or describe the way they handle a given task *versus* indirect methods which infer what the expert must have known from responses given by the expert, or from observations of the expert; examples are multidimensional scaling (see, for example, Olson and Reuter, 1987) and repertory grid analysis (see, for example, Hart, 1986; Boose, 1986; and Neale, 1988).
3 Techniques which are focused on a specific task or problem domain *versus* those that are generalized; examples of techniques for specific tasks or problems are STUDENT, an acquisition technique tied closely to the statistical consulting domain (Gale, 1987) and FIS, tied closely to the fault diagnosis domain (DeJong, 1986); an example of a general or domain independent technique is KITTEN, which is based on repertory grid analysis (Shaw and Gaines, 1987).
4 Authentic sources, i.e., human expert(s) *versus* reconstructed sources (textbooks, data bases, theories, etc.), of expertize.

The second set of literature describes KA as a phase within the larger expert system development life cycle (see Guida and Tasso, 1989, for an overview). The various approaches to characterize this life-cycle can be categorized into two general types: the technology-led or "rapid-prototyping", and human-centred or "non-rapid-prototyping" approaches. A more refined framework to classify the approaches has been suggested by posing the query of "what should drive the elicitation" (Cordingley, 1989, p98); there are four different ways to answer the question:

1 System or implementation driven: implementation refers to a particular knowledge representation formalism, or a chosen shell which constrains the information elicited. This approach is popularly referred to as "rapid-prototyping". It describes knowledge elicitation to take place with the Knowledge Engineer (KE) interviewing the expert (E) and encoding the verbal knowledge into rules. Then through iterative sessions, the E comments on and helps to debug and augment the knowledge base until the system performs satisfactorily. Buchanan, *et al.* (1983) supported this notion of technology-led or "rapid-prototyping". They outlined the major

stages of knowledge acquisition as identification, conceptualization, formalization, implementation and testing, and considered it advisable to move directly to the formalization and implementation stages as soon as key concepts and relations had been identified:

> "It will be tempting to try to analyze the problem correctly and completely before implementing a trial system, but experience has shown that to be a mistake. Once the key concepts and relations are written down, much can be gained from formalizing them and working toward an initial implementation." (Buchanan, 1983, p 144).

Recently, Weitzel and Kerschberg (1989) proposed a knowledge-based system life-cycle approach which they labelled as a prototyping methodology for knowledge-based systems. While this approach indicates a series of steps through which the developers have to work, the model specifically allows for the reactivation of any phase in order to refine a previous activity based on feedback received on the current prototype.

2 System and domain driven approach: both domain knowledge and implementation considerations influence the knowledge engineering process. Hart (1986) supported this approach. Similarly, Fellers (1987) viewed knowledge engineering to consist of two phases: knowledge acquisition and expert system construction. The former process involves "one or more knowledge engineers interacting with one or more domain experts (DEs)", and the objective of this interaction is to develop a "shared representation" or model of the DE's problem-solving process. Once the shared or external representation has been developed, the system construction phase begins. The KE takes the external representation and maps that to the actual knowledge base (KB). The KA process then continues with the expert interacting with the prototypical expert system to develop the internal representation or mental model of the system. The objective is to increase the congruence between this internal representation of the system and the mental model of the expert (Fellers, 1987).

3 Development process driven approach: in this approach, the requirements of the development process provides the framework for determining the relevant deliverables. Grover (1983) described the KA cycle as comprising the three stages of domain definition, fundamental knowledge formulation and basal knowledge consolidation.

4 Model-driven approach: this approach is exemplified by the KADS methodology, which suggests a structured approach to KA. It is based on a life cycle model in which knowledge analysis precedes the design and implementation of the knowledge based system (see Breuker and Wielinga, 1987*a*,*b*, and 1989 for detail).

Since the objective of this research study is to investigate the process of KA as it is implemented by practitioners in their development of expert systems, both sets of literature are relevant, and suggest notions about the phenomenon which will be useful points of comparison with observations derived from the case studies. The comparison between practice and theoretical prescriptions is discussed in detail in section 5.

To provide some initial focus to this empirical investigation, we used the model of interaction in the KA process (Gaines, 1987*b*) to identify the actors and processes involved. The study focuses on two processes, primarily on the knowledge acquisition, and secondarily on the expert system construction process; and the actors involved include the expert, the knowledge engineer and the user. The variables identified to be of interest include:

1 Problem domain.
2 Expert system.
3 Actors involved:
   (a) expert;
   (b) KE;
   (c) potential users.
4 Knowledge acquisition process.
5 Expert system construction process.
6 Organization setting.

3. Research strategy and design

3.1 Research Strategy

A case research strategy was adopted for this study for the following reasons. It has been suggested that case research is appropriate for investigations into areas "where the existing body of knowledge is insufficient to permit the posing of causal questions, and when a phenomenon cannot be studied outside the context in which it naturally occurs" (Bonoma, 1985, p 207). Furthermore, we wanted to study the development activity as a process over time in its natural context to understand the influence of the variables identified as potentially important by Gaines (see section 2 above). The phenomenon of interest, i.e., the acquisition process, takes place in organizations that currently develop such systems; an inquiry into the undertakings would shed light upon the question of how the systems are produced. In this sense, our study is different from Littman's (1987). He conducted a laboratory study to identify the knowledge and skills utilized by knowledge engineers in capturing an expert's underlying model of the domain. His study focused on one individual only, namely the knowledge engineer, but the interaction of the knowledge engineer with the expert was not examined. In fact, the knowledge engineer had to infer from a transcript of a dialogue between a tutor and a student what the expertize of the tutor was. In contrast, our study examined the interactions and the dialogue between the individual involved in the design as the knowledge acquisition process was taking place.

The four stages in the case research process are referred to as drift, design, prediction, and disconfirmation, which reflect a theory/data/theory revision cycle (Bonoma, 1985). The drift stage is of primary interest because our investigation addresses this stage. Bonoma explains it as follows:

> "This stage consists of the investigator's attempts to learn the concepts, locale, and jargon of the phenomenon as it occurs 'in the field', and to begin preliminary integration from literature, *a priori* notions about the phenomenon's operation, and critical components of practice as observed."

In other words, this is a "soaking in" stage in which the context as well as the content of the case are observed. A case constructed in this phase serves both as a problem for individual analysis and as a stimulus for further thinking [Bonoma, 1985, p.205]. Our study addresses this "soaking in" stage, it aims to understand and describe not necessarily to quantify or enumerate. In drawing lessons from the cases which form the basis for future hypothesis generation, the study also attempts to deal with the second phase of "design".

3.2 Research Design

Instead of examining a single case, this enquiry empirically investigates three expert system development projects. We also wished to have an arm's-length, objective examination of what was taking place in contrast to studies which report the experiences of those who developed the system themselves. Adopting an empirical approach to observational methodology (Weick, 1968, p 401), a questionnaire was compiled which addressed the variables identified in section 2. It was used when the first author observed the sessions for the first two cases, and interviewed the participants for the third (see the Appendix for the questionnaire).

3.2.1. Unit of analysis and site selection

The unit of analysis consisted of the knowledge acquisition process for an individual expert system. In practice, however, the clear distinction between this and the implementation or system construction process often cannot be made. Therefore, whenever necessary, the unit of analysis is extended to include the system development process as well.

The sites were selected based on two criteria:

1 An expert system was being developed.
2 Permission to conduct the study was procured.

The three cases investigated include the Remoteness Advisor developed at the Faculty of Law of the University of British Columbia, the Test Desk Decision Support Tool developed at Microtel Learning Services Centre, and the Wastewater Treatment Expert System developed at the Civil Engineering Department of the University of British Columbia, all located in Vancouver, Canada.

3.2.2. Data Collection

The study aimed to track the knowledge acquisition and development processes as they occurred. This allows the first author to study the phenomenon within its natural setting, and also ensures maximum currency of data. Bonoma (1985, p 200) noted the major weakness in case research to be low data integrity. To compensate for this weakness, triangulation of data, a technique which uses different types of data to capture a single event, was adopted. Hence, this investigation adopted multiple data collection methods, which include the following:

1 Direct (unobtrusive) observation of the knowledge acquisition and development processes between the expert and knowledge engineer (Webb *et al.*, 1966);
2 Taped protocols of the sessions;
3 Retrospective interviews with experts and KEs;
4 Diary or journalistic record of events kept by the KE;
5 Documentary evidence or deliverables produced during the process.

Since subjects' co-operation was essential, the above list was consistently presented to the KEs in the three projects. In two of the three cases, they agreed to co-operate on all methods except keeping a diary. In the third case, since the authors did not learn of the project until it was under way, only *post facto* interviews were conducted with the expert and KE using the developed questionnaire, and some written documentation from the KE was also obtained.

4 Description of cases

4.1 The Remoteness Advisor

The Remoteness Advisor is an expert system developed by a faculty member (the expert) and a doctoral student at the Faculty of Law of the University of British Columbia. The expert has written a book (Smith, 1984), in which he laid out a coherent theory of negligence developed after 23 years of teaching, research and writing in this area. The KE was pursuing a doctorate in the interdisciplinary program of law and computer science. He is a lawyer who had practised for ten years before returning to school. The two had previously co-operated in developing another expert system in the domain of law called the Nervous Shock Advisor.

4.1.1. Negligence: definitions

A central issue in negligence is the problem of *remoteness*, i.e., how far down the chain of causality could a person be liable for a negligent act so that a damage would be recoverable. Negligence usually involves the creation of risk of harm to someone. The injuries of the primary subject of the risk are generally not very remote; the remoteness issue arises when other persons suffer a loss as a result of the negligence to the *primary subject* (Smith, 1984, p 127).

Foreseeability of harm is a necessary condition for culpability, and consequently for the liability derived. The obvious case of foreseeability and culpability is where the harm caused is the very intent of the actor, in which case the actor is held liable for the intentional causing of harm (Smith, 1984, p 130). For any negligent act there are certain harmful effects which are foreseeable; when they materialize as a result of an action, the remoteness issue is *not* involved. There are other effects, however, which could not be said to be reasonably foreseeable, and it is in regard to these that the issue of remoteness of damage may be raised (Smith, 1984, p 132).

4.1.2. Objectives and nature of the task
The objectives of the expert system (ES) are:

1 To establish whether negligence and remoteness is involved, and if so, to categorize the negligent act that has led to the damage;
2 Once it is determined that remoteness is involved, to retrieve all prior court cases which could be used as precedents in court.

The major objective of developing the ES is the second one. The extent to which the ES is able to identify precedent cases that are as similar (or dissimilar) as possible to the one that is to be argued in court is one of the key measures of its value to the lawyers. However, similarity in the domain appears to be an elusive concept. The task thus becomes to find a scheme, involving both a structure and a procedure, to classify a given situation in order to be able to compare it to prior cases. As such, the task is unstructured, and the scheme to be chosen to a large extent depends on the expert's judgement as to what is a good match between the given situation and prior cases, and what factors, concrete and conceptual, have to be taken into account to make such a determination.

4.1.3. Data collection methods
KA for the Remoteness Advisor took place at the expert's office. There were a total of six KA sessions; each took 1.5 hours on average. The sessions took place between October 1987 and February 1988. The first author was present to observe, as well as tape-record, the discussion between the expert and the KE. For the one session where she was absent, the KE taped the session.

Before or after the sessions, and at occasions when the observer found it difficult to follow the legal jargon, independent or joint interviews with the expert and the KE took place.

4.1.4. Knowledge acquisition
It is difficult to describe the individual steps that the expert and KE went through during the process of knowledge acquisition. This in essence was a creative process of idea generation. Thus the discovery process was not entirely transparent to the independent observer, nor to the two participants who were asked to describe their thinking process. The description of the KA process will thus emphasize mostly the external aspects as captured by the observer. There are enough data that could be identified through such means which show how this KA project is different from the other two cases examined, as will be described later.

At the first KA session two distinct topics were discussed; the structure of the system and the presentation format. The structure of the expert system was different in the minds of the expert and the KE. They had obviously spent time pondering the issue before coming to the meeting. It was equally clear that they approached it differently. The KE suggested posing a question at the start to ask if there had been damage suffered as a result of a negligent act. The expert, however, preferred to ask questions which would guide the user to think in concrete terms such as the activity which caused the damage, the damage itself, and the risk that justified the possibility of positively concluding a causal relationship between the activity and the damage.

The structure of the system was a large categorization scheme, hence a major portion of the time during the KA session was spent on clarifying the activity, damage and risk involved in a case. A rudimentary theory of negligence which was available in the expert's book was used as a starting point. Other sources of knowledge used in the KA process were the previous case histories, that were used both inductively to develop the categorization scheme and later to test its completeness, and the knowledge and experience the expert and KE possessed as lawyers.

The role of the KE at this stage was twofold:

1 To ask the expert how specific instances should be categorized, and pose questions to clarify the expert's definitions;

2 To suggest the addition of some new categories to the basic scheme described in the expert's book. For example, a new category of radioactivity was added to the list of causes of damage which previously only had an item on inflammable substances.

Certain factors made this KA process more likely to succeed. First, the KE and the expert had worked on a similar project together. It was evident that they had a *close working relationship*. Often, they finished off sentences for each other. The KE agreed in the interview that since he was a lawyer by profession, he and the expert had a *shared vocabulary* which was a considerable advantage for him as a KE. This was especially true in the domain of law, where a deep understanding of concepts, as well as a real understanding of how practising lawyers function, was necessary for the KE to follow the expert's reasoning. In spite of this, conflicts of opinion arose. On issues dealing with the domain of negligence, the KE deferred to the expert's opinions, but the KE's suggestions and queries often clarified the matter in the expert's mind. The KE tried to *pose hypothetical cases to test the expert's categorization scheme*, and he used current news topics such as AIDS, the space shuttle disaster, and dangers of pit bull terriers. He needed to attach the expert's ideas to *concrete examples* in order to make the ideas meaningful to him.

Although the purpose of the KA session was explicitly to clarify on a *conceptual level* the domain of negligence, much discussion was on presentation details. The expert had definite ideas on what the resulting system would be like, and often suggested to the KE design details on presentation formats. He always had in his mind a picture of how the interaction with the system will appear to the user. The KE, on the other hand, admitted he had no idea what the system would look like.

The KE further pointed out *it was impossible to create rules at this stage*. It should be reemphasized that in their previous efforts together in developing the Nervous Shock Advisor, the structure was discovered after the rule base was complete and a working prototype available. As a result, fundamental structural changes were made during coding which led to friction between the expert and KE who was reluctant to make changes after investing a substantial amount of time in program development. This previous experience and the lessons learned reinforced the importance of diagramming the structure before beginning to code. It established an understanding about the need for a lengthy and detailed KA process, and illustrated a weakness of the rapid prototyping approach.

The second KA session for the Negligence Advisor took place three weeks after the first session. Once again, their discussions focused on the issues of the structure of categories, and presentation format.

The expert and KE continued to work on the underlying structure of the expert system. In the previous KA session, they already noted down a number of causes. They started by reviewing the categories and discussing various detailed issues such as ordering of the categories, specialization within general headings, and groupings of various items.

Presentation of the categories to users was also addressed in the session. There were discussions about how to pose questions to the users so that they would be able to reduce the number of queries to be answered. The expert pondered the possibility of starting with generally applicable categories and "get(ting) them out of the way." The KE encouraged him in this endeavour by posing more scenarios that involved combinations of categories such as a car crashing into explosives, or a train carrying radioactive materials.

The most interesting point about this KA session was the gradual emergence of the underlying "skeleton" for the expert system as the discussion developed. It seemed that the overall structure emerged despite the explicit focus of the discussion on specific details within the framework. They used the approach of bypassing the difficult and abstract and always trying to pose concrete questions to the user.

At the post-session interview, the expert further elaborated upon this method by emphasizing they left legal rules, philosophy, and anything disputational out of the system. In doing so, they avoided the risk of users not understanding the carefully formulated and articulated questions on legal philosophy. Rather, the system posed questions generated by the philosophy pertaining to the concrete situations. The KE pointed out that users, being lawyers, were not interested in theory but only in cases to give them some bases for making successful arguments in court.

The essence of their approach was to ascertain facts in cases by relating them to precedents which had similar facts. To do that one had to reason by analogy to find alike or relevant cases. This implied the need for a criterion of relevance. The manner in which a case was broken down to its basic elements was important. This would determine how a given situation would be compared to prior cases. If these basic elements were too general then too many cases would be retrieved; if they were too specific then some valid instances might be missed. Ultimately, the empirical test of success and validity was whether the user obtained the relevant cases and nothing was missing. The domain of law is distinctive in that the set of court cases which embodies all the facts to be considered is readily available. In other words, the scope of the cases relevant to the domain is clearly defined.

The third KA session for the Remoteness Advisor took place on December 7 1987. The KE started the session by reviewing what was accomplished the last time they met. They concluded there were two basic categories:

1 Special types of people, which included rescuer, thin skull (those prone to injury), nervous shock and medical complications; and
2 Special types of activities, which included weaponry, firearms, toxic substances, animals and motor vehicles etc.

Based on the conviction that the categorization was derived from the cases, the KE had examined the cases again to see how well they fitted into this scheme. According to the KE, without analysing all the cases there could be no certainty that all the key elements in the domain had been identified. He came away with proposals about how to modify the categories at hand into more generic ones and how to subdivide some categories further.

At this point, the expert proposed a second categorization scheme. Instead of looking at the nature of the vehicle, a second categorization scheme that cut across all vehicles was needed. The expert felt that this second level categorization, which was independent of the kind of vehicle involved, was important and needed to be clarified on top of the first level categorization that the KE had formulated. He suggested that they re-examine all transportation cases, and identify the plaintiff's relationship to the risk in order to seek the patterns among them. (As explained by the KE, it is naïve to interpret a case in only one way, thus multiple representation schemes of the same cases would provide a more realistic simulation of legal reasoning.)

At this session it is interesting to note the use of cases (1) to measure comprehensiveness of the scheme developed to date, (2) to generate new categories and refine others, and (3) to develop a more abstract concept of remoteness.

The fourth KA session for the Remoteness Advisor occurred one week later. The expert and the KE continued their development of the underlying structure of the system. The KE pointed out that in their previous work the structure had been derived from the exercise of developing the system, but the expert felt that if they could arrive at a structural diagram before developing the system, the KE could work backwards from the structure and expand it into a massive set of rules and questions. In this way, the process of KA would be accomplished in half the time it took for the Nervous Shock Advisor. Or alternatively, if they could arrive at the structure and work in both directions, the availability of the structure would help clarify to the students working on the knowledge base of cases what information to look for. This again is an advantage which the rapid-prototyping approach does not enjoy.

In this session, the expert and the KE tried to arrive at the second level categorization that the expert thought was necessary in the previous KA session. The first author observed that this attempt had not been completely successful as the second level categorization was by no means obvious, and there were times during the session when the participants seemed stuck. Considering the expert's objective of arriving at the underlying structure of the system before the system was developed, it seemed he was embarking upon an innovative effort to short-circuit their original KA

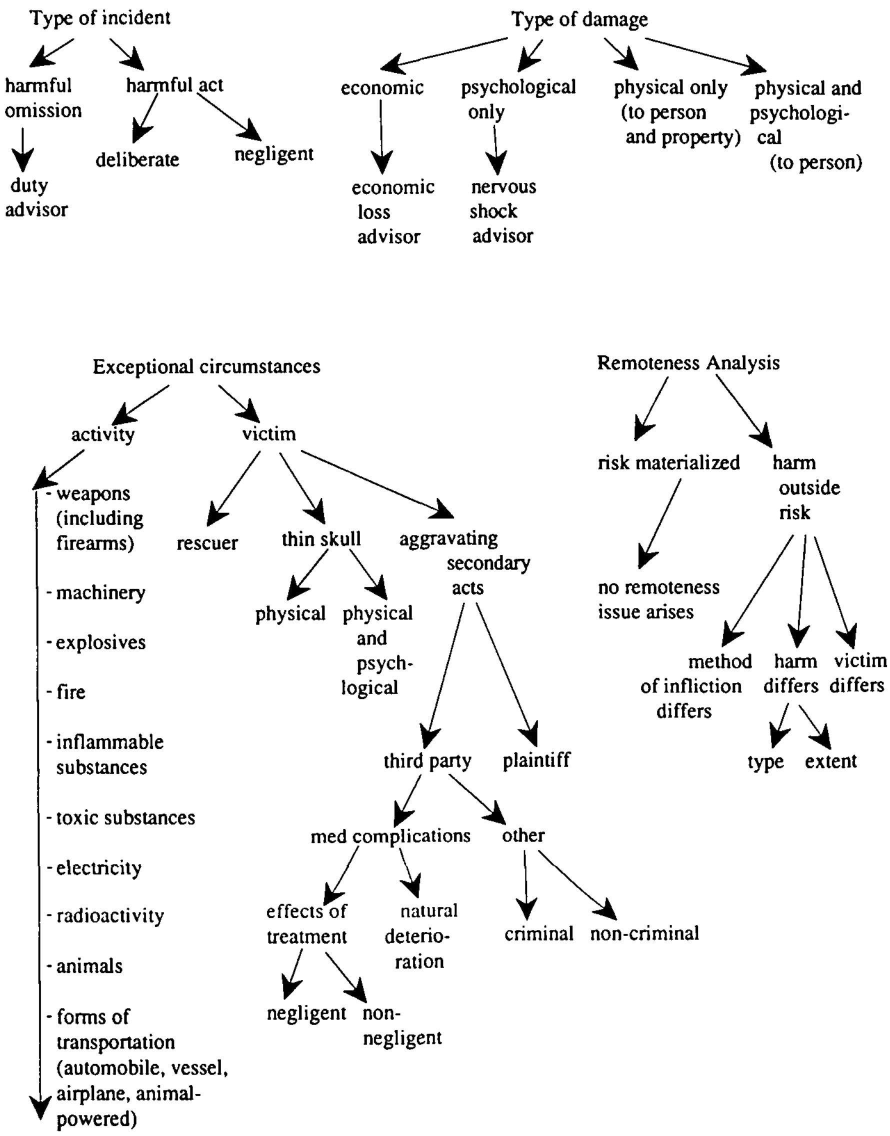

Figure 1 Remoteness Advisor structural diagram.

efforts. It remained to be seen if a structure could be derived in this manner without what the KE felt was the necessary, but painful, experience of fumbling through the expert system development process.

During the Christmas holidays, the KE constructed a schematic diagram to categorize all negligence cases; the structure that emerged was in fact the skeleton of the system. The expert had approved of the structure, and in early January and February of 1988, they met to further refine the structure (see Figure 1).

They examined individual branches of the diagram and posed actual or hypothetical cases to each other, and deliberated on how a particular case would travel through the branches, whether it would fall into one category *versus* another. The objective was apparently twofold: first, to label the categories generally enough to capture all relevant cases, and second, if the category was found to be inadequate, insert another one. The KE was concerned that the terminology should be general enough to cover all possible cases of that type, but also that it should not be overly-academic for lawyers.

Both thought in terms of the overall expert system design, and often adopted the perspective of the user travelling down the system. The expert often took the initiative in assuming the role of the user.

At the end of the session in February, they agreed that the session was a milestone for the end of the KA phase and the beginning of the prototype stage. The KE decided to test the control flow by implementing a small version in M1.

One last session was spent on refining the structure. Then, the KE worked independently from the structural chart and derived the list of variables which would be used to characterize the cases. Based on the structural chart, he developed a prototype of the Remoteness Advisor on M1, which was shown to the first author in April 1988.

4.2 The Test Desk Decision Support Tool

The Test Desk Decision Support Tool is an expert system developed at the Microtel Learning Services Centre. The KE, a senior instructional designer at the centre, acted as manager of the expert system project. He has a Master's degree in Instructional Systems Technology, and has worked for over ten years as a courseware designer. The chief expert on the project is a senior instructor at the British Columbia Telephone (BC Tel) Education Centre, with over 30 years of experience. The chief expert worked closely with the KE over various phases of the development process; but as discussed later, five other experts also participated in the project.

4.2.1. Test desk operation

The expert system serves as a job aid for the test desk person (tester). Presently, three to four such persons are involved in processing the trouble reports produced by the telephone testing equipment system which performs daily routine checks on the telephone lines. When the reports arrive at the Trouble Analysis Centre (TAC), the TAC supervisor assigns a tester to investigate the trouble reported. The tester gives his analysis report to the clerk, who in turn makes up a dispatch to send out a TAC isolator to the site. The isolator calls from the site and the test desk person sends out test signals to the isolator. They interactively test the line until a diagnosis can be made and the problem remedied.

If a wrong diagnosis has been made then the time and cost incurred is substantial. An expert system could potentially incorporate the knowledge that the best test desk persons have and aid an average tester to more accurately diagnose the problem. Not only is the time and money saved in making accurate diagnosis important to the company, but speedy line repairs improve customer satisfaction.

Despite the wealth of information a tester has to analyse, there are only 14 types of trouble and six types of dispatches to the TAC isolators which are identified and discussed in the manuals. Given a trouble type, only a finite number of procedures could reasonably apply, hence the domain is definable in this respect.

Instead of relying on a single expert, the KE planned to draw from the aggregate of all sources, which included the manuals for the two testing systems (Badger and 4tel), descriptions of office procedures, human informants such as supervisors of test desk persons, expert test desk operators, as well as "non-expert" desk persons (to see why they fail in their work). Furthermore, approximately one third of the information for the expert system was already resident in a simulation program developed previously, and could be transferred directly to the new system. The KE was aware of this because he had participated in the simulation software development.

The expert system was being developed simultaneously as the materials for a course on the testing equipment was being updated. The expert system was to serve as a job aid and training tool. Its intended users included testers in the operational environment and trainees in the classroom at the Education Centre. The chief expert would be teaching the course and his primary interest was in revising the materials. The KE helped him with the revision while at the same time extracting materials from it for building the expert system.

4.2.2 Data collection methods
Organizational funding and availability of personnel were obstacles with which the KE had to contend from day one of the project. Several times, development was suspended due to these problems, while other times, progress was made between one day and the next. The first author tracked the development by staying in touch with the KE by telephone and attended sessions personally about once a week from November of 1987 to early February of 1988. The data were collected either over the telephone from conversations with the KE and sometimes with the chief expert, or through direct observations when the first author attended the sessions and tape-recorded them.

4.2.3 Knowledge acquisition
Knowledge acquisition for the test desk decision support system was carried out under the rubric of "task analysis."

The objective was to decompose the job of the test desk person to its component duties, tasks, subtasks and prerequisite knowledge. An example of a task decomposition is shown as follows:

1 Duty: Identify and Diagnose Faults
Task:
1.1 read trouble report (TAS)
1.2 interpret trouble report
1.3 perform basic line test
1.4 interpret basic line test
1.5 determine type of trouble
1.7 determine what test to perform
1.8 perform 4tel tests
1.0 interpret test result
1.10 route/write off

The first phase of the analysis involved verifying and updating the materials in the two testing equipment system task analysis manuals (the manuals for the Badger and the 4tel systems were compiled in 1981 and 1986, respectively).

The KE and expert cross referenced the two task analysis manuals to see how the job had changed. They updated the material and decided what was to be included in the expert system. Although some of the materials were seven years old, and the tools used today were different, most of the knowledge was still valid. However, there was considerable *variety across different locations*; 55% of the company sites used the 4tel testing equipment and the rest had either the Badger system or none at all. Since the sites varied in size and complexity, the office procedures were different. To standardize the procedures among the diverse methods adopted was difficult.

Typically, the expert would read through the 1986 manual, checking each step with his notes and the 1981 manual. Sometimes he raised a point, or the KE asked him if some step was indeed currently implemented. Either they arrived at a conclusion or one of them made a note to call somebody for clarification.

At the end of the initial phase, the KE and the expert discussed the timing of presenting what they were doing to the staff personnel. Such a presentation was necessary both to obtain the staff's verification that indeed the KE and expert had captured the correct and current method of conducting the tasks, and to reconcile the different methods that existed in the field.

In the week of November 30 1987, the KE and chief expert conducted group KA sessions for five days with four experts brought in from various areas in British Columbia. The purpose of the meeting was primarily to gather and verify information to be included in the course and the expert system. Again, they based their discussions on the two existing manuals compiled in 1981 and 1986. The KE and the chief expert had begun this work and now the four local experts were brought in to continue the verification process. This process involved a verification and update of the two manuals. One person from the group would read from a copy of the manual and as he was doing so, the others made suggestions on appropriate changes and clarifications.

In this process, substantial discussion arose due to diversity in equipment and procedures among the various centres. *Hence considerable time was spent when the participants exchanged information on procedural differences.* Even though the chief expert acted as the arbitrator in the event of disagreements, he often was unable to declare a company standard because procedures vary from region to region.

Despite these differences, it was agreed the basic duty of the test desk person was to "test reported trouble" and the six major dispatches involved to achieve this were generic and common across all sites, regardless of equipment and tools. Similarly, there existed only 14 trouble types possible. Focusing on the generic, they updated the materials in the manuals.

After verifying the information on the documents to be accurate and current, they proceeded to integrate the two documents by listing each duty and its subtasks on the board. Discussion arose over the proper decomposition of the tasks and whether some tasks should be deleted. The integration that emerged represented what the five experts consented was the best way of conducting the duties. The new arrangement was noted down on the master manual and sent to the typist so that a new version of the manual could be produced.

The sessions were also used as an information dissemination forum so the test desk personnel from different districts had a chance to exchange insights and identify problem areas. The expert noted down these and subsequently wrote letters to the staff personnel to request further investigation into identified problem areas. He also made recommendations that certain procedures could be passed onto the test desk personnel in other districts.

The next stage of the KA process had the experts discuss how they would handle different trouble types. The task involved correctly interpreting the trouble reports, accessing the necessary databases and analysing the data to arrive at a diagnosis, which would point to a type of trouble, and then to send out the appropriate dispatch. Again working from the flowcharts in the manuals, the experts described the steps necessary to verify that the problem was indeed caused by a particular type of trouble.

The difficulty in the group sessions was reconciling the differences across the sites. Interaction among the four experts during flowcharting was noticeably more intense than before. The method to handle different trouble types varied from person to person, and the discussion at times became heated as they could not agree on one best way to handle a trouble type. When heated arguments erupted over fine points in the flowchart, they were either resolved by the group or by the chief expert's promise of writing memos to check with some other persons in the company. The KE was skilful as the mediator of the meeting and in making sure everyone had a chance to participate.

On the last day of the group meetings, the KE explained the top level design of the expert system to the group. Working out the procedure to be encoded in the knowledge base and clarifying part of the static information had been the objective for the last five days. The KE explained they aimed to build a prototype based on the "no-dial-tone" type of trouble by January or February of 1988, and then expand upon the knowledge base. By the time the prototype would be completed they would meet again to evaluate it.

In mid-December, the KE and chief expert met for the first time since the group sessions two weeks earlier. Their current objective was to concentrate on the flowchart for the trouble type of no-dial-tone, which covered 40% of all trouble reports and whose diagnosis was considered relatively straightforward.

4.2.4 Development of the rule base

The KE, together with an outside consultant, co-operated in developing an initial prototype. They worked from the flowchart used to diagnose the no-dial-tone type of troubles. The consultant anticipated problems with writing the rules as she was unfamiliar with the test desk problem, and the KE promised to help her in writing the rule base.

In mid-January, the KE familiarized himself with the shell and then wrote the rules for the no-dial-tone flowchart up to the test/no test decision point. Subsequently, the consultant identified the syntactically incorrect rules. As the KE explained the meaning of the syntactically incorrect rules to the consultant and the latter edited them, they were simultaneously discussing the structure of the knowledge base. They iteratively refined the rule base by putting down a few rules, and then backtracked and ensured all the antecedents were covered, and added a rule to explain them if necessary.

Another test desk person, who was a senior technician with seven years test desk and two years 4tel maintenance experience, was assigned to the project to substitute for the chief expert who had become unavailable because of involvement in other company projects. The new expert and the KE focused on refining and adding new rules and modifying the wording of the explanatory text. The consultant worked on correcting the syntax of the rules and ensuring the control structure worked. She probed for clarifications on how something was done, and the KE and test desk person would check the hardcopy rule base in response to her queries.

The test desk person was asked to examine the behaviour of the system. As he tried different options, he compared the system's responses with his experience to verify that it reacted as he would at the test desk. In his judgement, the system contained all the rules necessary to handle no-dial-tone type troubles up to the test/no test decision point.

Since the company demonstration of the system was scheduled for early February, they spent a substantial amount of time discussing how to prepare for it. The consultant pointed out the interface was "terrible". Presentation of the line of reasoning was difficult for users to understand and she planned to "pretty it up." The prototype in general was not easy to use, and did not have a good display interface.

In the beginning of February, the chief expert tested the prototype. He assumed the user's role and selected different options to test the system's logic. He detected a logic error in the program, some non-standard acronyms and typos, and that two fields were missing in the database. Recognizing the prototype only dealt with simple problems within one trouble type, the chief expert suggested incorporating a more involved trouble type, and providing window access to more than one report screen from the database whenever necessary. The KE acknowledged these areas to be important, but decided to postpone incorporation of more rules to handle other trouble types. Presently, they would concentrate on improving the interface to the report screens in preparation for the demonstration.

4.2.5 Epilogue

The KE reported that in early February of 1988, a presentation was made to the staff group at BC Tel, which included Management Information Systems (MIS) personnel, service centre managers and union representatives. The MIS people were most receptive and interested in using the technology in other applications. They tentatively agreed to commit people and money to support the project. A second demonstration was given to the service office managers from all areas in the province.

The consensus so far was to use the expert system for internal training and in the courses. But district managers showed reservation in using it at the operational level. They were unsure how it would fit into current and future test desk operations, and deferred decision to the Staff Manager of Operational Support Services, who had championed another innovation before. In the expert's opinion, the expert system had been developed using a "grass-root" or "bottom-up" approach, and the KE had been the champion behind the project. The time had come for top management to accept the technology; such acceptance is usually a prerequisite for successful implementation (Prerau, 1989).

In mid-February, the KE reported that the presentations to senior management of the company received a "favourable" reaction The MIS group showed keen interest in expanding the expert system technology to other applications, and he was making plans to bring the system to the other branch offices for feedback.

4.3 The Waste Treatment Expert System

The *Was*te *T*reatment *E*xpert *S*ystem (WASTES) was developed by a Master's student in environmental engineering at the department of Civil Engineering of the University of British Columbia.

The KE had taken a number of courses on wastewater treatment, and was familiar with the domain. The experts who contributed knowledge for the system included his thesis supervisor (the chief domain expert) and three plant operators. The thesis supervisor is a recognized authority in the design of wastewater treatment plants. The three operators were chosen based on their extensive experience in the field

4.3.1 Data collection methods

Data on the development process was obtained primarily from *post facto* interviews with the KE and the chief domain expert. Concurrent observation of the process was impossible because the authors did not learn of the project until six months after its initiation, when only two days of coding remained to complete the first prototype. The following narrative description had been compiled from the interviews and from a draft of the thesis (Chilibeck, 1988).

4.3.2 Wastewater treatment operations

The system aims to aid the operation of conventional secondary activated sludge wastewater treatment facilities and diagnose problems that may arise. Treatment facilities are called preliminary, primary, secondary and tertiary, depending on the level of waste water treatment provided by the combination of unit operations. A treatment plant consists of an interconnected string of unit operations, which are physical, chemical and biological processes that remove waste materials from the water. A plant operator controls conditions such as temperature, air, nutrients and hydraulic flows in the pipes in order to maintain the system to facilitate or retard bacteria growth. Other variables that can be controlled include age and type of bacteria.

The secondary treatment facility was chosen as the model for the system because the largest proportion of plants in use in North America belongs to this category, and developing a system for this level of treatment would describe preliminary, primary and secondary treatment operations. The treatment of the municipal wastewater is usually carried out by the *activated sludge process* in which micro-organisms metabolize the waste material. The activated sludge process is suitable for expert system technology because it requires substantial operator knowledge and input for process control as compared to other secondary processes.

Traditional process control of waste treatment plants has been done algorithmically by manual or automatic control systems. WASTES however, uses heuristic rules based on observational data on the color and condition of the mixed liquor and aeration tanks, and its microscopic inhabitants. The KE noted that expert system technology improved upon numerical or algorithmic process control because it was more adaptable to conditions unforeseen by the operator, and was often cheaper. Numerical control programs required installation of digital meters at critical sites at the plant so that data on the conditions could be fed into the program. This required considerable investment and may not be suitable for plant operators who are generally not technically sophisticated. Their knowledge of the plant is mostly gained through experience, and their method of control is, according to the KE, "an art more than a science." The expert likewise suggested that numerical equations were inadequate for process control. The experience of operators at the local plants served as a more reliable basis for control than equations.

For the KE, motivation for developing the system also came from the need to fill the gap between research theory and treatment plant operations. The experience that operators possess fills this gap and encoding it in heuristic rules is one means to capture this valuable knowledge. A new operator learns his trade by example and from the knowledge of the experienced operators. The system is therefore useful as an operational as well as teaching aid. Anticipated users of the system include experts such as waste treatment plant design engineers as well as non-experts like municipal engineers, who are unfamiliar with the domain, and students.

4.3.3 Knowledge acquisition

There were three sources of knowledge for the system, and their relative importance was as follows:

1 50% from published works of process control of waste treatment plants, which supply primarily diagnostic information;
2 25% from the chief domain expert, the thesis supervisor; and
3 25% from the plant operators with whom the KE had short conversations over the telephone, primarily when he needed to clarify process control information.

The written sources of information consisted of texts specifically written for plant operations and process control, articles from magazines, and research publications from journals. The Environmental Protection Agency has published troubleshooting guides and control manuals for municipal treatment. Additional troubleshooting information was gained from the Water Pollution Control Federation publications on design and operation of treatment facilities.

From the literature, the KE obtained information which he represented in logic trees. Then he asked the expert-operators for verification of the information presented in the logic trees. Since every plant was different and no model-plant existed, the analysis presented in the written works often did not coincide with reality. For example, the KE presented to the expert-operators three different ways of controlling a process derived from three different schools of thought. He asked them which method they adopted at their plants, and the operators chose what the KE considered to be the simplest one from among the three. The expert then modified the logic tree accordingly, and the corresponding rules were incorporated into the knowledge base.

4.3.4 Development of rule base

After the initial phase of acquiring knowledge, the KE consolidated the information he needed for the expert system by identifying the following:

1 Unit processes and operations to be included; and
2 Common plant problems.

The result of this consolidation process was a logic structure or flowchart for each unit operation, relating all possible problems with the parameters to be examined, which determine what precise problems existed, and solutions for them. The KE then proceeded to convert the flowcharts, starting with the smallest and simple ones, into rules, which were subsequently coded and debugged for logic and syntax errors. A series of mini-knowledge bases were developed in this manner, one for each unit process. Then they were linked to the menu structure in the main module of the program and to the knowledge base. Rule bases on problems of floating materials and odour, which affected the entire plant, were also included.

As the knowledge base was being developed, the expert's contribution of knowledge played a significant role. Some processes were not well documented in the literature, thus he *supplied the necessary clarifications.*

According to the KE, knowledge acquisition sessions with the expert were highly informal and usually lasted half an hour. For simple problems, the KE described the scenario to the expert, and through discussion arrived at a conclusion. For more complex problems, the KE demonstrated the system to the expert, who assumed the user's role and suggested changes when necessary. The KE reported that he consciously maintained the expert's interest by periodically showing him the system. Their interaction was facilitated by a shared vocabulary. Hence, continual feedback was obtained from the expert primarily by posing scenarios and letting him see how the system functioned. On several occasions, the expert also took the system home over the weekend when he reviewed it as a user to check if all relevant variables had been considered.

In situations when the three operators did not agree among themselves, the KE relied on the expert's judgement. With years of experience as a plant design engineer, the expert was able to resolve differences or attribute them to circumstantial variances. If operators disagreed on the optimal procedures because of differences in plant size and conditions, they might all be correct for their own plants. In such cases, the system had to accommodate all situations.

While the operators were knowledgeable about the plant operations, the expert had an overall appreciation for the entire operation of the treatment facilities as well as how the environment, such as the quality of the incoming sewage, impacted the plant. At an interview, the KE stressed the importance of the operators' input for "a good operator can rectify a lot of design errors in a treatment plant; a good designer still needs a good operator at the other end or you wouldn't get results."

The expert added that while he was helping the KE with his system, he was simultaneously working closely with a plant operator to improve the waste water treatment operations at a local plant. Discussions with the KE often clarified his own understanding of the problem at that plant.

In the expert's opinion, WASTES is usable but not foolproof. It is complete in that it has addressed all the important variables that should be considered in a conventional secondary treatment plant. But for a less conventional facility that involves processes not adequately explained by current, state-of-the-art knowledge in environmental engineering, the system would have problems. More basic research in the field of process control in wastewater treatment plants is required before a "complete" system can be developed.

4.3.5 Epilogue

In early April of 1988, the KE gave a demonstration of WASTES to his fellow students at the Civil Engineering Department. At the time of writing, he planned to send the system out for evaluation by the field operators.

5 General observations and limitations of the study

The study has reported on three expert system development projects. A number of observations drawn from the inquiry and some limitations of the study will be discussed in this section.

5.2.1 General observations

1 While the "rapid-prototyping" approach is a well-established technique in building expert systems (Gaines, 1987*a*; Buchanan *et al.*, 1983), the notion of "rapidness" is hard to define precisely. Buchanan *et al.*, for example, suggested implementation should start once "a few key concepts and relations" are noted down. Similarly, Naumann and Jenkins (1982) and Jenkins (1983) emphasized the importance of constructing a working prototype within a short time. Observations of the law and test desk cases suggest that while prototyping is a valid technique in building expert systems, implementation *should not* begin until the shared representation of the domain is reasonably understood by both the KE and the E.

Jenkins has stated that prototyping for information systems design rests on the premise that development of complete and correct information requirement specifications prior to implementation is both technically and behaviourally infeasible (Jenkins, 1980). For expert systems, instead of specifying the information requirements, expertize is to be represented. The mapping between expertize and the captured model requires a series of modelling processes, and prototyping is appropriate for ES construction precisely because it is a vehicle of communication between the KE and the E (Jorgensen, 1984).

A clear definition of prototyping is often lacking in the literature. In the present context, it refers to a system that is to be expanded, supplemented, or supplanted (Jenkins 1980), and not to a system that is "created as quickly and as cheaply as possible to test the feasibility of a new concept or approach" (Asner *et al.*, 1987). In fact, observations from all three cases suggest that the latter notion of prototyping is to be avoided, and implementation should begin only when the domain addressed is substantially clarified.

In the case of the Remoteness Advisor, the expert and knowledge engineer spent seven sessions formalizing the structure of the domain (and the system) before the KE developed a prototype. At that point, both were reasonably confident that the structure would remain stable. Speaking from experience, the expert and KE stressed the importance of coding only *after* the structure of the domain had been formulated. Prototyping prematurely had resulted in substantial changes in the code in the prototype system during their earlier work together. This caused conflict between the expert and the KE, because the latter found it difficult to discard his initial prototype. This observation is consistent with concerns about the rapid-prototyping approach raised in Diaper (1989) and Brueker and Wielinga (1987*b*).

Similarly in the case of the Test Desk Decision Support Tool, the project team spent two months updating and verifying the knowledge in the manuals before they took the conceptual design of the expert system to the software developer. This was done even though the cost devoted to clarifying the knowledge was substantial, since the four test desk persons had to be released from their jobs for a week and brought to the headquarters in Vancouver for the group sessions.

From these two cases where prototyping was not immediately attempted, a number of observations can be made. In the Test Desk case, there was a reasonable amount of structure in the task domain as codified partially in procedure documents and expressed as the scheme of task decomposition in the task analysis manuals. Furthermore, the number of trouble types and resulting dispatches were reasonably small. Therefore, it seemed reasonable to work around this structure and gather data from experts to update the current knowledge base before a prototype was developed. In the Remoteness Advisor case, while the structure was elusive in the beginning and therefore prototyping appeared to be a logical alternative, a prototype was not immediately attempted. This was probably due to the disagreement that arose between the expert and KE, when in their previous work a substantial amount of coding was discarded due to discovery of a new structure after a prototype was implemented. That episode strongly influenced the behaviour in this case.

2 The KA techniques that were observed in the three cases include the following:

(a) Interview in the form of question-answer between the expert and KE was observed in all three cases. Typically, unstructured interviews predominate at the early and structured ones at the later KE sessions.
(b) In all three cases, extensive use was made of non-authentic expert sources, i.e., books and documentation.
(c) Scenario-elicitation, where the KE posed a scenario and the expert commented on it was observed in the two engineering cases. In the design of the law system, cases were used to elicit or induce the underlying structure and to test the emerging structure.
(d) Task analysis by decomposing the domain into goals and subgoals was used in the test desk case.
(e) Group discussions among the experts were observed in the test desk case.
(f) Prototype or system refinement was observed in the two engineering cases.

The use of scenario-elicitation supports the idea that experts find it easier to demonstrate their expertize in a particular situation rather than formulate it explicitly (Witten *et al.*, 1988). Such techniques are by no means the only appropriate ones: prototype or system refinement would be another way to elicit scenarios. In general, no one technique is either sufficient or best to conduct KA. Rather, KA is composed of a series of techniques applied at different phases of the expert system development life cycle, each technique serving a different purpose. It is also interesting to note that there is usually a wealth of existing information in written documents that could be used before the input of human experts is sought. However, written documents belong to what has been referred to as "reconstructed methods of reasoning", and may not reflect authentic expertize (Johnson, 1983).

3 The nature of the task domain affects the KA process by influencing the type of experts to be consulted. In the two engineering cases, since the tasks involved variations in local conditions and could not be adequately modelled as a single universal process, multiple experts were required (Henderson, 1987). Only the Remoteness Advisor case falls within the traditional mode of eliciting from one expert. In this case, the domain addressed is more subjective than the other two, and the expert was one of the few proponents of a "deep structure" interpretation of the law of negligence.

The domain is also important in another respect. Both the Test Desk and Wastewater Treatment cases involved diagnostic tasks, while the law system addresses a classification problem (Hayes-Roth *et al.*, 1983). In the former instance, the number of parameters to be examined varies with the local conditions, while in the latter the number of relevant facts are dictated by the set of precedence-setting court cases for this domain.

4 It has been observed that the stages in the expert system life cycle in fact overlapped (Buchanan *et al.*, 1983). The inquiry supports this finding. Despite the KE's efforts to keep the discussion to a conceptual level, the expert in the Remoteness Advisor case frequently considered implementation issues during the KA sessions. The expert and KE of the Test Desk Decision Support Tool recalled screens or data that could be used from the simulation program. A segregation of the stages of conceptualization and implementation rests on the assumption that the domain expert is ignorant about information technology. In reality, however, experts who are willing to participate in expert system projects are individuals highly interested in computer technology and eager to be involved in developing the systems. In the case of the Wastewater Treatment Expert System, the expert claimed to be non-computer-oriented, and the scheme of separating the conceptualization and implementation stages prevailed.

5 Close rapport between the expert and KE directly and positively influences the quality of the conceptual model produced. The law case is distinctive in that personal friendship and professional collegiality existed between the expert and KE, and the conceptual structure that resulted from the KA sessions was a product of the intellectual as well as personal dynamics that took place (Smith, 1988).

The law and civil engineering cases are similar in that both KE and experts were knowledgeable in the domains. The KE in the law case is a lawyer by training, and the KE for WASTES a graduate student in environmental studies. Hence, these KEs shared the experts' language as well as knowledge. In the law case, for example, the KE was often able to cite relevant cases to test the expert's categorization schemes; a non-lawyer would not have had such a facility.

6 The Remoteness Advisor and the Wastewater Treatment Expert System were developed in an academic setting while the Test Desk Decision Support Tool was produced in a commercial environment. The pressure of building a prototype to win organizational and managerial support applied in the latter but not in the former two cases. These organizational realities influence the role of the KE; he might sometimes have to act as a system champion and consensus builder. It also influences the speed with which a working system has to be developed with an eye on resource constraints, such as the availability and assignment of funding and personnel. This in turn might lead to the use of techniques that will expedite the development of a prototype in a narrow piece of the domain to show feasibility and to assure continuity of organizational commitment.

7 Handling the communication, conflict reduction, and consensus reaching activities among multiple experts is a major time consuming task in expert systems development. Two models of such activities have emerged in the cases studied. In the test desk case the experts interacted with each other and with the KE, at times and in situations where their opinions were in disagreement. While the existence of a primary expert having final authority is desirable (Prerau, 1989), the test desk case indicates that the chief expert is not always knowledgeable about differences

between local conditions. In the waste treatment case the KE interacted with three local experts and a chief expert, but the experts did not communicate with each other. An emerging information technology, Group Decision and Communication Support Systems, would be an ideal one to facilitate the management of KA from multiple experts, especially in conditions similar to the one encountered in the test desk system where consensus building in constructing the system was very important (Kulikowski, 1989). Such systems use information technology to support communication and collaboration between group members, idea generation, idea organization, planning, creating, and the elicitation of knowledge to create expert systems (Stefik *et al.*, 1987; Nunamaker *et al.*, 1988). The use of such systems would significantly increase the efficiency, and possibly decrease the cost, of expert system development with multiple experts.

5.2 Observations on the data collecting methodology

Conducting longitudinal observation of multiple cases requires a research team (Yin, 1984). This theoretical prescription, however, is impossible to practice in the present study because the subjects involved did not agree to having alternative observers. Having familiarized themselves with the first author and having expended the time and effort to acquaint her with the basic jargon, subjects were understandably reluctant to reinvest their energies. While cases do provide the benefits of current data and observation facilitates understanding, the demands of keeping track of multiple cases over a period of almost a year were onerous for one person.

Observation of the KA sessions was a valuable means to obtain data. Maintaining rapport and developing trust with the subjects were essential so that they did not feel uncomfortable about being watched.

Taped protocols of the sessions were useful if the quality of the tapes could be assured. This proved difficult, especially since it was impossible to attend every single session and subjects who took over the responsibility might have been less careful. Group sessions were difficult to tape, and the recordings were hard to transcribe as multiple voices might overlap. In the law case, however, since there were only two participants in a reasonably small room, the tapes were invaluable when legal jargon became difficult to comprehend during the sessions. A complete transcript of each tape proved useful because relevant information on development of the knowledge structure occurred at sometimes unexpected points during a session.

Retrospective interviews were useful for clarification purposes, and less so as a primary data gathering tool. In the case on the Wastewater Treatment Expert System, for example, the first author had no means of validating responses obtained. The interviewees might have succumbed to retrospective biases.

Not all subjects responded favourably to keeping a diary, and in the three cases described, it was not adopted. In another two cases not reported in this paper because the systems were abandoned halfway through the projects for various reasons, the KEs involved agreed to keep journalistic accounts of the processes, which proved extremely informative in tracking how they approached their tasks.

5.3 Limitations of the study

1 This study depended heavily on observational data. Despite the use of deliverables, taped protocols and *post facto* interviews with subjects to triangulate the collected data, reactive bias due to presence of the first author was minimized but not avoided. Since the first author observed the KA sessions in the legal and the test desk cases for a period of over six months, reactive bias introduced probably decreased as the subjects became more accustomed to being watched. Another threat to internal validity came from the first author's record and interpretation of data, which was controlled by showing the reports on each case to the subjects involved for verification.

2 External validity is constrained by the fact that only three cases were studied. Attempts to generalize the lessons to other development efforts can be made only with due consideration to the domain addressed and the human actors involved.
3 The emphasis of this study has been on the KA process. In all three cases, only a demonstration-oriented prototype had been developed by the time observation terminated. Hence the entire development life cycle, i.e., stages subsequent to the prototype development phase, was not documented.
4 The sample is biased in that the reported projects all succeeded in producing working systems. Cases examined did not include failures.

6 Concluding comments

This study has investigated the process of KA using the case research strategy. As explained in section 3.1, the objective of the drift stage of case research is to *understand* through data collection. Observations from this study can serve as basis for induction in the next stage of the inquiry.

The data gathered in the cases can be used in different ways in the design stage of the case research process. One possibility is to more formally identify and investigate the variables of the KA phenomenon initially addressed during the drift stage of the research. Dhaliwal and Benbasat (1990) have proposed a framework to investigate the efficacy of knowledge acquisition which includes the following factors:

1 Dependent variables: these are the variables that measure the quality of the expert system developed and the efficiency with which the KA process was conducted.
2 Independent variables: these are the KA tools and techniques which will potentially influence the dependent variables described above.
3 Moderator variables: these variables will influence the selection of the KA tools and techniques as well as the relationship between the dependent variables and the KA techniques chosen. They include:

   (a) the human actors (expert and KE);
   (b) the characteristics of the problem domain;
   (c) the system development life cycle methodology chosen; and
   (d) the organizational environment.

The observations of this study apply primarily to the independent, and secondarily to the moderator, variables. It has confirmed that the techniques for KA and the process of conducting KA are strongly affected by the background and experiences of the experts and KE, by the nature of the task domain, by the number of experts involved in the development, and by organizational factors such as availability of resources. While this study has by no means answered many of the questions associated with KA, it has attempted to add to the literature, which is at an early stage in terms of empirical research, by objectively tracing the process of KA for developing three expert systems.

Acknowledgements

This work has been supported by operating grant OGP2421 from the Natural Sciences and Engineering Research Council of Canada. The authors would like to express their sincere gratitude for the co-operation of the participants in the three cases. In addition, we would like to thank Dr. John Fox and the two anonymous referees for their constructive comments on an earlier draft of the paper.

References

Abrett, G and Burstein, MH, 1987, "The KREME knowledge editing environment" *International Journal of Man-Machine Studies* **27** (2) 103–126.

Alexander, JH, Freiling, MJ, Shulman, SJ, Rehfuss, S and Messich, SL, 1987, "Ontological analysis, an ongoing experiment" *International Journal of Man-Machine Studies* **26** (4) 473–485.

Anjewierden, A, 1987, "The KADS system" *Proceedings of the First Knowledge Acquisition Workshop* Reading University, UK, (September).

Asner, M and King, A, 1987, "Prototyping: a low-risk approach to developing complex systems" unpublished paper from Gellman Hayward & Partners Ltd.

Bell, J and Hardiman, RJ, 1989, "The Third role—the naturalistic knowledge engineer" In: Diaper, D, ed., *Knowledge Engineering: Principles, Techniques, and Applications* Ellis Horwood Ltd.

Bonoma, TV, 1985, "Case research in Marketing: opportunities, problems, and a process" *Journal of Marketing Research* **22** (May) 199–208.

Boose, J, 1986, *Expertise Transfer for Expert System Design* Elsevier.

Boose, JH, 1989, "A survey of Knowledge acquisition techniques and tools", *Knowledge Acquisition* **1** (1) 3–38.

Breuker, J and Wielinga, B, 1987*a*, "Knowledge acquisition as modelling expertise: the KADS methodology" *Proceedings of the First European Workshop on Knowledge Acquisition for Knowledge-Based Systems* Reading University, UK, (September).

Breuker, J and Wielinga, B, 1987*b*, "Use of models in the interpretation of verbal data" In: Kidd, A, ed., *Knowledge Acquisition for Expert Systems: a Practical Handbook* Plenum Press.

Breuker, J and Wielinga, B, 1989, "Models of expertise in knowledge acquisition In: Guida, G and Tasso, C, eds., *Topics in Expert System Design: Methodologies and Tools* 265–295. Elsevier Science Publishing Co.

Buchanan, BG, Barstow, D, Bechtel, R, Bennett, J, Clancy, W, Kulikowski, C, Mitchell, T and Waterman, DA, 1983, "Constructing an expert system" In: Hayes-Roth, F, Waterman, DA, and Lenat, DB, eds., *Building Expert Systems* Addison-Wesley.

Burton, AM and Shadbolt, NR, 1987, "Knowledge engineering", Technical Report 87-2-1, Artificial Intelligence Group, Department of Technology, University of Nottingham.

Chilibeck, BM, 1988, "Operation and diagnostics of wastewater treatment facilities using an expert system", unpublished Master's thesis, Department of Civil Engineering, University of British Columbia, Canada.

Cordingly, ES, 1989, "Knowledge elicitation techniques for knowledge based systems" In: Diaper, D, ed., *Knowledge Elicitation: Principles, Techniques and Applications* Ellis Horwood Ltd.

DeJong, K, 1986, "Knowledge acquisition for fault isolation expert systems" *Proceedings of the First Knowledge Acquisition for Knowledge-Based Systems Workshop*, 20.0–20.8, Banff, Canada, (November).

Dhaliwal, JS and Benbasat, I, 1990, "A framework for the comparative evaluation of knowledge acquisition tools and techniques" *Knowledge Acquisition* **2** (2) 145–66.

Diaper, D, 1989, "Designing expert systems: from Dan to Beersheba" In: Diaper, D, ed., *Knowledge Elicitation: Principles, Techniques, and Applications*. Ellis Horwood Ltd.

Fellers, JW, 1987, "Key factors in knowledge acquisition" *Computer Personnel* **11** (1) 10–21.

Gaines, BR, 1987*a*, "Rapid prototyping for expert systems" In: Oliff, M, ed., *Intelligent Manufacturing: Proceedings of the First International Conference on Expert Systems and the Leading Edge in Production Planning and Control* 45–73 Benjamin Cummins.

Gaines, BR, 1987*b*, "Advanced expert system support environments" *Proceedings of the Second Knowledge Acquisition for Knowledge-Based Systems Workshop* Banff, Canada, (November).

Gale, WA, 1987, "Knowledge-based knowledge acquisition for statistical consulting system" *International Journal of Man-Machine Studies* **26** 55–64.

Greenwell, M, 1988, *Knowledge Engineering for Expert Systems* Ellis Horwood Ltd.

Grover, MD, 1983, "A pragmatic knowledge acquisition methodology" *Proceedings of the Eighth International Joint Conference on Artificial Intelligence* Morgan Kaufmann Publishers, Inc.

Guida, G and Tasso, C, 1989, "Building expert systems: from life cycle to development methodology" In: Guida, G and Tasso C, eds., *Topics in Expert System Design: Methodologies and Tools* North Holland.

Hart, A, 1986, *Knowledge Acquisition for Expert Systems* Kogan Page.

Hayes-Roth, F, Waterman, DA and Lenat, DB, 1983, *Building Expert Systems* Addison-Wesley.

Henderson, HC, 1987, "Finding synergy between decision support systems and expert systems research" *Decision Sciences* **18** (3) 333–349.

Jenkins, AM, 1980, "The prototype model as a management information systems design technique" Discussion Paper #1663, School of Business, Indiana University.

Jenkins, AM, 1983, "Prototyping: a methodology for the design and development of application systems" Discussion paper #227, School of Business, Indiana University.

Johnson, PE, 1983, "What kind of expert should a system be?" *The Journal of Medicine and Philosophy* **8** 77–97.

Jorgensen, AH, 1984, "On the psychology of prototyping" In: Budde, R, Kuhlenkamp, K and Mathiassen, L, eds., *Approaches to Prototyping* Springer-Verlag.

Kitto, CM, 1988, "Progress in automated knowledge acquisition tools: how close are we to replacing the knowledge engineer?" *Proceedings of the 3rd AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop* Banff, Canada (November).

Kulikowski, CA, 1989, "Knowledge base design and construction: from prototyping to refinement" In: Guida, G and Tasso, C, eds., *Topics in Expert System Design: Methodologies and Tools* North Holland.

Littman, DC, 1987, "Modelling human expertise in knowledge engineering: some preliminary observations" *International Journal of Man-Machine Studies* **26** 81–92.

Naumann, JD and Jenkins, AM, 1982, "Prototyping: the new paradigm for systems development" *Management Information Systems Quarterly* **6** (3) 29–44.

Neale, IM, 1988, "First generation expert systems: a review of knowledge acquisition methodologies" *Knowledge Engineering Review* **3** (2) 105–145.

Nunamaker, JF et al., 1988, "Knowledge-based systems support for information centers" *Journal of Management Information Systems* **5** (1) 4–24.

Olson, JR and Reuter, HH, 1987, "Extracting expertise from experts: methods for knowledge acquisition" *Expert Systems* **4** (3) 152–168.

Prerau, DS, 1989, "Choosing an expert system domain" In: Guida, G and Tasso, C, eds., *Topics in Expert System Design: Methodologies and Tools* North Holland.

Shaw, MLG and Gaines, BR, 1987, "Techniques for knowledge acquisition and transfer" *International Journal of Man-Machine Studies* **27** 251–280.

Smith, JC, 1984, *Liability and Negligence* Carswell Legal Publications.

Smith JC, 1988, Professor J C Smith emphasized this point in personal communications with the first author and C Deedman on August 19 1988.

Stefik, M, Foster, G, Bobrow, DG, Kahn, K, Lanning, S and Suchman, L, 1987, "Beyond the chalkboard: computer support for collaboration and problem solving meetings" *Communications of the ACM* **30** (1) 32–47.

Sviokla, JJ, 1986, "Planpower, Xcon, and Mudman: An in-depth analysis into three commercial expert systems in use" unpublished doctoral thesis, Graduate School of Business, Harvard University, Cambridge, USA.

Watkins, P and O'Leary DE, 1987, "Knowledge acquisition for small scale expert systems from consulting experts: a field study approach" *Proceedings of the First European Workshop on Knowledge Acquisition for Knowledge-Based Systems* Reading University, UK (September).

Webb, EJ, Campbell, DT, Schwartz, RD and Sechrest, L, 1966, *Unobtrusive Measure: Nonreactive Research in the Social Sciences* Rand McNally.

Weick, KE, 1968, "Systematic observational methods" In: Lindzen, G and Aronson, E, eds., *The Handbook of Social Psychology Volume 2* Addison-Wesley.

Weitzel, JR and Kerschberg, L, 1989, "Developing knowledge-based systems: reorganising the system development life cycle" *Communications of the ACM* **32** (4). April.

Welbank, M, 1987, "Knowledge acquisition: a survey and British Telecom experience" *Proceedings of the First European Workshop on Knowledge Acquisition for Knowledge-Based Systems* Reading University, UK (September).

Witten, IH and MacDonald, B, 1988, "Using concept learning for knowledge acquisition" *International Journal of Man-Machine Studies* **29** (2) 171–196.

Yin, RK, 1984 *Case Study Research, Design, and Methods* Sage Publications.

Appendix: questionnaire

An Investigation of Knowledge
Acquisition Process in Expert System Development

Prepared by
Christine Chan
Faculty of Commerce and Business Administration
University of British Colombia

What do we want to learn?

I Environment

Problem

1 What is the problem domain of the expert system?
2 Why is this problem suitable for expert system development?

Actors

3 Why do you use one, a few, or many experts?
4 What are the criteria for choosing the expert(s)?

Setting

5 Describe the organization or its subunit, for which the expert system is being developed (e.g. size, employees, nature of business).
6 Where would the development take place? (at a company? at BCIT? or other?)

II Process

1 Date and time of each event or activity related to ES development. What happened (e.g., make a phone call to cross check a piece of info)? Why did you do that? Who were involved?
2 Milestones (or critical points) in the series of activities or events.
3 Iterations (looping back to previous events) involved?
4 How is knowledge acquired?
   (1) What is the method of knowledge elicitation?
   (2) Do you present the expert with problems to solve?
   (3) Do you ask directly the rules s/he uses?
   (4) Do you observe how s/he goes about solving the problem? Where does that take place?
   (5) Do you lay any ground rules for interacting with the expert? (e.g., set up common glossary, fix time for one session, etc)
5 How is knowledge tested?
   (1) Do you develop hypotheses of how the expert reasons and test them?
   (2) Do you compare what the expert says with how s/he actually solves a problem?
   (3) Do you compare one expert's method of solving a problem with other experts' approaches?
   (4) Do you develop prototypes and then refine them?
6 How is knowledge represented?

(1) How do you organize the knowledge?
(2) How is knowledge acquired translated into machine representation?

7 Other Considerations
(1) Do you try to maintain commitment and interest level of the expert?
(2) How do you do that?

III Outcomes

1 How do you know if the system is complete?
2 How do you evaluate its performance?
3 Do you anticipate problems with user acceptance? And if so, any strategies to ensure success?

How are data to be collected?

1 Keep log on the following:
(a) date and time of each event or activity related to expert system development.
(i) What happened?
(ii) Who was involved?
(iii) Why did that happen or why did you do that?
(b) Any milestones and iterations in the journal of events recorded in (a).
(c) When applicable, please note down any response or thoughts relevant to the previous questions.

2 Tape-record every interaction with the expert(s).
3 Researcher to be present at some critical sessions; e.g., when you interview experts or at prototype demonstrations.
4 Summary of data sources:
(1) Diary,
(2) Direct observation,
(3) Taped protocol,
(4) Retrospective interviews with knowledge engineer.