A PRAGMATIC KNOWLEDGE ACQUISITION METHODOLOGY*

Mark D. Grover

TRW Defense Systems Group
One Space Park
Redondo Beach, California 90278 USA

ABSTRACT

A formal, but pragmatic, method of recording and organizing human expertise into a knowledge-based system is presented. Practical considerations and methods which increase system validity while minimizing demands on human domain specialists are explored.

The methodology concentrates on domain definition (background knowledge, references, situations, and procedures), on fundamental knowledge formulation (elementary rules, beliefs, and expectations), and on basal knowledge consolidation (review and correction cycles). Experience gained from testing and from expert feedback is described.

I BACKGROUND

Until recently, Artificial Intelligence researchers have developed knowledge-based systems (KBS) for problems in which the human expert was constantly available and involved. However, the very scarcity of the expertise to be modelled may be the motivation for system development. The expertise may be heuristic and dispersed among several individuals while a single expert may be difficult to locate, let alone interview.

Traditionally, the knowledge acquisition (KA) phase of a new application has taken one of two approaches. In the first, an existing model well-suited to the new domain is used to develop a set of axioms and rules, i.e., the knowledge base. Such is the case with the PUFF program (Kunz, 1978) under the EMYCIN paradigm (van Melle, 1980). A second KA method is the historically successful team approach where a domain expert and a knowledge engineer are locked away for months or years and return with a model and a computer program which are comparable in performance to human specialists. This method has produced several remarkable programs such as PROSPECTOR (Duda, 1978), INTERNIST (Miller, 1982), and R1 (McDermott, 1980).

Other experimental approaches such as discovery systems (Lenat, 1982) may be practical eventually. Meanwhile, the knowledge engineer must still resolve the problem of limited availability of experts in disciplines where the expert is unique or indispensible and cannot be spared from the day-to-day task. For example, there are many human advisors in political or military domains producing situation assessment and warning reports (Ben-Bassat, 1982) who are also performing other critical functions. They often cannot dedicate months to develop an expert system which could be used to aid in the synthesis of the huge amounts of data which decision-makers must consider in order to warn of attack.

Access to these experts may be restricted due to reasons of organization, geography, or security. Additionally, the knowledge in political and military domains is often sketchy, diversified, distributed, and highly transitory. The more efficiently a knowledge engineer spends the available time, the more valid the model produced. It is reasonable to work on the assumption that any methodology which improves knowledge acquisition productivity in these extreme cases will also be helpful in less stringent cases. Additionally, expert system verification (Suwa, 1982) is a necessary component of formal knowledge acquisition and any limitations placed on the availability of human expertise places corresponding limitations on the scope and validity of the models and programs under development.

We must explore domains which break many of the traditional precepts upon which cautious knowledge engineering practice has relied. The expert may not be an available, authoritative individual, but instead a collection of part-time advisors aiding the engineer in the creation (rather than the recording) of consolidated expertise. We are in the early stages of approaching this problem, but present techniques and results are relevant to traditional knowledge acquisition tasks. These techniques may also be applied in the more general case of specifying solutions to large-scale software engineering tasks which use heuristic approaches and algorithms.

II KNOWLEDGE ACQUISITION CYCLE

Although generally undocumented, many knowledge acquisition techniques are intuitive, common practices. A significant innovation is the production of a KA Document Series. The establishment of this documentation would be a partial substitute for the expert and would provide system designers and users with a common medium for communication and reference.

The suggested knowledge acquisition methodology for a new problem domain takes place in three phases: domain definition, fundamental knowledge formulation, and basal knowledge consolidation (Figure 1).

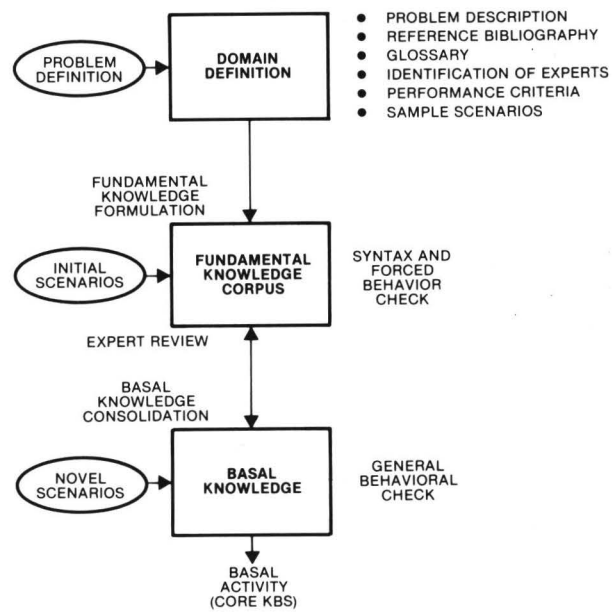

Figure 1. Knowledge Acquisition Cycle

The documentation series contains the results of the three phases and allows users, experts, and system designers to possess a consistent, organized, up-to-date set of written human experience on which to base the KBS.

*This work was sponsored by 1982 internal TRW Research and Development funds.

III DOMAIN DEFINITION

After the problem is defined by the user and requirements team based on accepted engineering practices, the first phase of knowledge acquistion is the careful understanding and recording of the domain. The goal is the production of a Domain Definition Handbook containing:

- General problem description,
- Bibliography of reference documents,
- Glossary of terms, acronyms, and symbols,
- Identification of authoritative "experts,"
- Definition of appropriate and realistic performance metrics, and
- Description of example reasoning scenarios.

The second item deserves comment. Source references are often sufficient to introduce the engineer to the domain; in particular, government sources have great volumes of useful documents, although they may not always be easy to obtain.

An informal experiment was performed to analyze effort expended on domain definition. A domain practitioner in Washington, D.C. and a knowledge engineer in Los Angeles communicated over a period of fourteen weeks by telephone and mail couriers only. The domain was a plausible but fictitious military situation assessment problem where the expert examines standardized reports of sensor signals and produces an interpretation of their meaning from heuristic knowledge. The effort expended, which included development of a typical scenario for later analysis but no rules, was

- 103 hours KA time (1/5 time)
- 47 hours expert time over 14 weeks (1/2 of KA time)
- 29 hours secretarial time (1/3 of KA time)
- KA labor: Interviewing 14%, Collating 14%, Analyzing 10%, Describing Orally 12%, Writing 15%, Background Study 21%, and Evaluating 14%

The most surprising result is the time spent in the telephone interview. Not so startling is the time spent in background study of documents and texts which could have been decreased by having the expert available to answer tutorial-level questions. This experiment produced 150 pages or approximately 40,000 words.

IV FUNDAMENTAL KNOWLEDGE FORMULATION

In the second phase of Knowledge Acquistion, scenarios are reviewed which have been selected by the expert and which fulfill five criteria of "fundamental" knowledge: the most nominal, the most expected, the most important, the most archetypal, and the best understood situations. This review forms a baseline for minimum performance, predictable testing and correction, and careful delineation of capabilities which can be expanded and subjected to experimentation. This fundamental knowledge baseline should include:

- An ontology of domain entities, object relationships (classes), and object descriptions;
- A selected lexicon (vernacular);
- A definition of input sources and formats;
- An initial state description including static "background" knowledge;
- A fundamental set of reasoning and analysis rules; and
- A list of human strategies (meta-rules) which may be considered by system designers for possible inclusion as control rules.

This corpus is written in English, rather than a formal language. Some of this knowledge will have previously been acquired during domain definition. Reasonable validity of this corpus can be tested only by presenting an implemented version of the knowledge with the scenarios from which it was acquired and producing the behavior expected by the original expert.

A. Interview Techniques For Rule Acquisition

Four different approaches for interviewing an expert during fundamental knowledge formulation were attempted in an unpressured setting; (1) forward scenario simulation (archetype acquisition), (2) goal decomposition ("twenty questions"), (3) procedural simulation (protocol analysis), and (4) pure reclassification (frame analysis). The first and last methods were the most useful and were later used in a highly time-restricted set of interviews.

In the *forward scenario simulation* approach the expert chooses a very elementary scenario and verbally "walks through" the reasoning necessary to reach a goal. This interview takes place in a laboratory setting, not an actual working environment of the expert. This approach, while apparently direct, has at least three difficulties:

1. Blazing a trail through uncharted territory involves terms and definitions whose details may not have been made clear in domain definition. Delays and confusion may result.
2. Rule fanout problems are recognized but must be ignored if one is to carry through with the scenario. A rule may become too specific or a reasoning path forgotten.
3. Distinction between the reasoning methods and the *job* of the reasoner may be confounded.

The second approach, *goal decomposition,* is the traditional problem reduction approach and is useful for enumerating goal states and describing general categories of goals; however, efforts to acquire detailed rules in reverse have not been successful. The interview may begin with "Suppose there is an X" but collapses at "What is prohibiting X from performing its mission?". Such rules are illustrative of purposes but do not produce detailed interpretations of situations and objects. Additionally, the fanout problem remains unsolved.

In the third approach, *procedural simulation,* a domain specialist is required to solve a specific problem in a normal manner and setting rather than to construct rules verbally. However, the interviewer must occasionally interrupt the natural reasoning process. Otherwise, the knowledge engineer may infer, rather than elicit, the actual reasoning (Simon, 1980). At the very least, this approach enables the engineer to observe and appreciate the expert's job, and may contribute to overall design improvements.

The fourth approach, *pure reclassification,* considers that much of the effort in a situation assessment problem domain is spent reclassifying observables into more specific objects and activities. The rules therefore will seek information to do that interpretation. It appears appropriate to guide the interview toward the construction of rules which fill in gaps in observations, similar to traditional backward chaining models.

Presuming that the problem is object-oriented and that these objects are well known to the expert, then the construction of reclassification rules may be useful. Although this construction also implies more complex rule control, it provides a means of sub-goal formation and allows work on many fronts at once. However, reasoning chains tend to be intricate and multi-leveled.

Knowledge engineers at TRW have commented that these hypothetical objects could be represented as "frames" with rules attempting to fill the slots in those frames. Eventually, enough information is gathered to conclude the existence of one or more specific objects and activities. On the other hand, simple object-attribute pairs or a relational model may be sufficient. Further, the term "frame" may be too inexact to be meaningful, and the simpler schemes of representation may be preferable. Once the domain is explored in sufficient detail, a more concise model can be constructed.

B. Follow-Up Study

A follow-up study was performed in which the engineer and expert from the domain definition experiment spent one week in the same location under intense interview conditions. Eighteen hours of actual interviews resulted in fifty-three preliminary English analysis rules for

the fictional military warning scenario and 127 background axioms and initial conditions. The refinement and implementation of these rules using a variety of computer languages is ongoing but incomplete. The following is one of the sample rules acquired.

IF SOME GUN OR MISSILE UNIT IS DETERMINED NEAR SOME LOCATION (LOC_1)

AND THERE IS SOME GUN OR MISSILE DEFENDABLE ITEM AT SOME LOCATION (LOC_2)

AND THE DISTANCE BETWEEN (LOC_1) AND (LOC_2) IS LESS THAN THE GUN OR MISSILE RANGE OF THAT UNIT

UNLESS THAT UNIT IS MOVING

THEN ASSERT THAT UNIT IS PROBABLY DEFENDING THAT ITEM

V BASAL KNOWLEDGE CONSOLIDATION

Knowledge acquisition is a continuous process and the third major step in this process is the "review and improve" cycle of basal knowledge consolidation (Figure 2). *Basal activity* may be defined in the same sense as in physiology: the lowest level of activity (system behavior) essential to the maintenance of vital functions. In a KBS, this refers to a core system in which all of the components in an operational system are performing, but without the breadth and depth of a rigorous environment. It must, however, meet the initial minimum performance standards set in the domain definition.

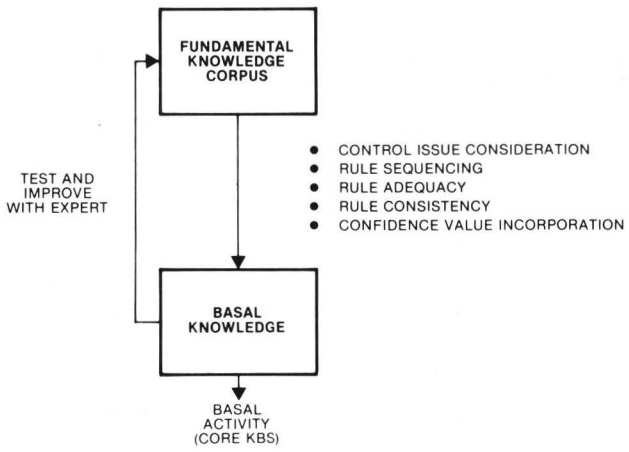

Figure 2. Basal Knowledge Consolidation

Basal knowledge, then, is that set of rules and definitions adequate to produce basal activity. The fundamental knowledge corpus is reviewed and enhanced by appropriate reconstruction of rules. Corroboration with additional experts may be attempted at this point. Confidence levels, if appropriate, are added. The expert is consulted in detail regarding strategy for the appropriate application and sequencing of reasoning rules; this strategy is used as "inspiration" for the selection of KBS control mechanisms. Resource bounds for programs are different than those for humans. Thus, control issues for attention management must be resolved on the basis of efficient use of computing resources not brain cells. Extensive behavior checks are performed both with the original expert and panels *a la* INTERNIST (Miller, 1982). This review process continues through the addition of novel scenarios until the system consistently behaves within the established general domain in a manner approved by the original experts. This core system can undergo further development as required.

VI COMMENTARY

Some comments gathered from the domain specialist (Snell, 1983) warrant special interest, e.g., avoid talking about performance expectations during the early phases; do not rush the expert; keep rules as self-evident as possible; expect ambiguity; use existing documentation; and let the expert do the talking.

The intent of this paper is to encourage others to discuss related experiences. This pragmatic paradigm for knowledge acquisition under limited access to expertise deserves greater exploration in order to minimize the waste associated with poor use of expert time and the frustrations of inadequate methodology.

These experimental results are somewhat anecdotal and further studies would be very useful. Examination of similar cases should aid in the development of important applications involving rare expertise in "fringe" domains, as well as improve more universal knowledge acquisition efforts. The methods described in this paper are being applied to contractual work (Grover, 1983; Goodman, 1983) as well as a new domain involving expert tuning of numerical model parameters of a complex system.

REFERENCES

[1] Kunz, J., et al. "A Physiological Rule-Based System for Interpreting Pulmonary Function Test Results." Report HPP-78-19. Computer Science Department, Stanford University, 1978.

[2] van Melle, W. "A Domain Independent System that Aids in Constructing Consultation Programs" Report STAN-CS-80-820. Computer Science Department, Stanford University. 1980.

[3] Duda, R. O., J. Gaschnig, P.E. Hart, K. Konolige, R. Reboh, P. Barrett, and J. Slocum "Development of the PROSPECTOR Consultation System for Mineral Exploration." Final Report, SRI Projects 5821 and 6415. SRI International, Menlo Park, California. 1978.

[4] Miller, R.A., H.E. Pople, and J.D. Myers "INTERNIST-I: An Experimental Computer-Based Diagnostic Consultant for General Internal Medicine." New England Journal of Medicine. August 19, 1982. 468-476.

[5] McDermott, J. "Rl: An Expert in the Computer System Domain." Proceedings of the National Conference on Artificial Intelligence. AAAI, 1980. pp. 269-271.

[6] Lenat, D.B. "Heuretics: Theoretical and Experimental Study of Heuristic Rules." Proceedings of the National Conference on Artificial Intelligence. AAAI, 1982. pp.159-163.

[7] Ben-Bassat, M. and A. Freedy "Knowledge Requirements and Management in Expert Decision Support Systems for (Military) Situation Assessment." IEEE Transactions on Systems, Man and Cybernetics. Vol. SMC-12:4. (1982) 479-490.

[8] Suwa, M, A.C. Scott, and E.H. Shortliffe "An Approach to Verify Completeness and Consistency in a Rule-Based Expert System." AI Magazine. 3:4. (1982) 16-21.

[9] Simon, H.A. and K.A. Ericsson "Verbal Reports as Data." Psychological Review 87:3 (1980) 215-251.

[10] Snell, D.J. "Observations on Being Knowledge Engineered." Unpublished TRW internal memo. 1982.

[11] Grover, M.D., H.M. Beebe, D.A. Brown, and D.J. Snell "Knowledge Engineering for the Adept Workstation." TRW Special Programs Report 38182.002. February 1983.

[12] Goodman, H.S. and D.A. Brown "Artificial Intelligence Applied to C^3I." Signal March 1983. 27-33.