Shaw, M. L. G. and Gaines, B. R. (1989). Extensions to knowledge acquisition methodologies and tools. In Proceedings of the Sixth IASTED International Conference Hamza, M. H., ed, pp. 76-79, Los Angeles.

Paper accepted for Sixth IASTED International Conference, Los Angeles, 14-16 December 1989.

Extensions to Knowledge Acquisition Methodologies and Tools

Mildred L. G. Shaw & Brian R. Gaines
Knowledge Science Institute
Department of Computer Science
University of Calgary
Calgary, Alberta
Canada T2N 1N4

ABSTRACT

Most researchers and practitioners agree that knowledge acquisition is the bottleneck in the task of developing knowledge-based systems. The basic model for knowledge engineering has been that the knowledge engineer mediates between the expert and knowledge base, eliciting knowledge from the expert, encoding it for the knowledge base, and refining it in collaboration with the expert to achieve acceptable performance. This basic model has been greatly extended by the introduction of interactive knowledge acquisition and encoding tools allowing the expert to transfer knowledge directly to the system without an intermediary. As these systems develop, they can take advantage of research and practice form other areas of computer science and other disciplines. Areas to be considered in the paper include: the automation of knowledge acquisition as a part of the evolution of information technology; the comparison of the use of terminology and conceptual structures from several experts; the increasing concern with and the development of methodologies for the validation of knowledge bases; the integration of machine learning and empirical induction techniques; the integration of text analysis, and hypermedia systems and tools; and the convergence with existing Computer Aided Software Engineering tools.

INTRODUCTION

The basic model for knowledge engineering has been that the knowledge engineer mediates between the expert and knowledge base, eliciting knowledge from the expert, encoding it for the knowledge base, and refining it in collaboration with the expert to achieve acceptable performance. Figure 1 shows this basic model with manual acquisition of knowledge from an expert followed by interactive application of the knowledge with multiple clients through an expert system shell.

This basic model has been greatly extended by the introduction of interactive knowledge acquisition and encoding tools allowing the expert to transfer knowledge directly to the system without an intermediary. Such tools can greatly reduce the need for the knowledge engineer to act as an intermediary, but, in most applications, they leave a substantial role for the knowledge engineer. As shown in Figure 2, knowledge engineers can:

- Advise the expert on the process of interactive knowledge elicitation
- Manage the interactive knowledge acquisition tools, setting them up appropriately
- Edit the unencoded knowledge base in collaboration with the expert
- Manage the knowledge encoding tools, setting them up appropriately
- Edit the encoded knowledge base in collaboration with the expert
- Validate the application of the knowledge base in collaboration with the expert
- Train the clients in the effective use of the knowledge base in collaboration with the expert by developing operational and training procedures

This use of interactive elicitation can be combined with manual elicitation and with the use of the interactive tools by the knowledge engineer rather than, or in addition to, the expert. The knowledge engineer can:

- Directly elicit knowledge from the expert
- Use the interaction elicitation tools to enter knowledge into the knowledge base

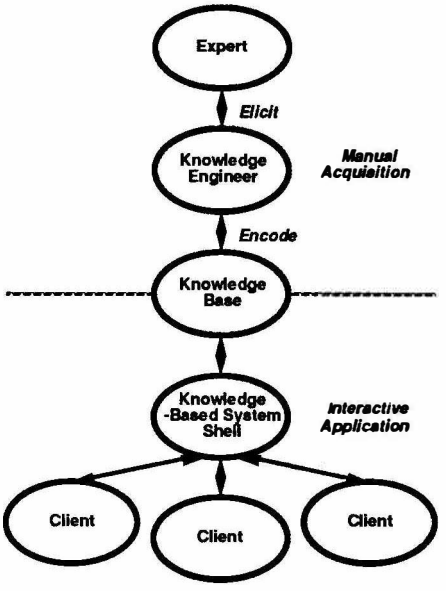


Fig.1 Basic model of manual knowledge acquisition

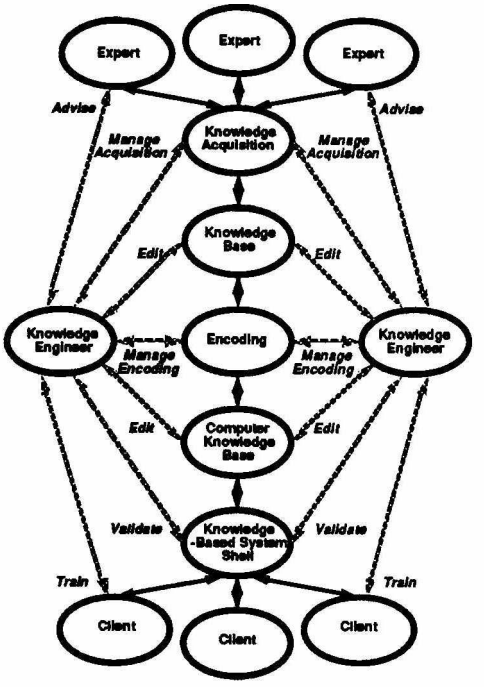


Fig.2 The knowledge engineers' roles in interactive knowledge acquisition

Figure 2 shows multiple knowledge engineers since the tasks above may require the effort of more than one person, and some specialization may be appropriate. Multiple experts are also shown since it is rare for one person to have all the knowledge required, and, even if this were so, comparative elicitation from multiple experts is itself a valuable knowledge elicitation technique.

TRENDS IN KNOWLEDGE ACQUISITION

The initial development of knowledge acquisition techniques in the 1970s was based on human interviewers debriefing experts about their skills. The development of computer-based tools supporting knowledge acquisition in the 1980s has been based on the computer-aided application of specific techniques such as construct and conceptual mapping, behavior modeling and text analysis. However, neither manual nor tool-based approaches have yet made a significant impact on the knowledge acquisition bottleneck which continues to impede the development and effective application of knowledge-based systems. Current research is attempting to establish a new paradigm for knowledge acquisition, one that recognizes the wide diversity of forms and sources of knowledge, its essential dynamism making it subject to contradiction and change, and the constructive nature of knowledge acquisition whereby it is often the knowledge acquisition process itself that creates knowledge.

Gaines (1988a) has suggested that second generation knowledge acquisition targeted on applications in the 1990s have to:

- Integrate a wide variety of weak methods encompassing the diversity of forms and sources of knowledge
- Integrate conflicting sources of knowledge to cope with essential contradictions that cannot be replaced by an artificial and misleading 'consensus'
- Integrate knowledge acquisition with system operation to facilitate a heavy activity of continuing system maintenance and upgrade as knowledge changes

- Integrate knowledge representation in the acquisition and performance systems, retaining the flexibility for new representations to be used that more accurately carry the knowledge

Thus, second generation developments have two objectives that may come into conflict:

- To integrate closely a variety of techniques and technologies to cope with the diversity of knowledge
- To provide flexibility within the integration framework for new techniques and technologies to be incorporated to cope with the dynamics of knowledge.

KNOWLEDGE AND INFORMATION TECHNOLOGY

To examine the technological support for second generation integrated knowledge acquisition systems it is necessary to have some model of the wide range of technologies underlying computing. The computer itself is only one component in an elaborate and highly differentiated infrastructure which has grown through a succession of generations of computers, each of which represents a major change in computing technology (Gaines & Shaw 1986). During the eight year span of each computing generation, revolutionary changes have taken place that correspond in magnitude to those taking place over some hundred years or more in other industries. The first generation commenced in 1948 and we are now in the sixth generation spanning 1988 through 1995.

The generation structure of information technology can itself be analyzed in terms of the knowledge flows involved. Like other basic technologies, the growth of computing has an infrastructure comprised of the envelope of learning curves in successive underlying technologies. Figure 3 shows this structure through the zeroth through fifth generations projected to the sixth and seventh. Each technology has a learning curve in which a *breakthrough* leads to a phase of: *replication* in which the results are duplicated; *empiricism* in which rules for design are derived from experience; *theory* in which generative principles are derived; *automation* in which the theory is made operational; and *maturity* when the knowledge becomes proven, widely available and widely used. The qualitative changes that have given computing its distinct character in each generation appear through the tiered succession of learning curves of higher level technologies based on the lower level developments.

The generations of knowledge acquisition systems are shown at the top of Figure 3—the current sixth generation of information technology is also marked. It coincides with the second generation of knowledge acquisition systems and this four generation gap indicates an analogy between the state of knowledge-based systems development now and that of data processing 32 years ago. We do indeed have a similar situation to the 1950s in computing when software development was new with problems of shortage of skills, unrealistic schedules and budget over-runs.

The detailed infrastructure for the development of the second generation of knowledge acquisition systems may also be seen in Figure 3. At the three lowest levels, the hardware machines architectures and languages are all mature technologies posing no problems. Human-computer interaction has entered its phase of automation with the increasing availability of machine-independent interactive systems such as X-windows. Knowledge-based systems have entered the theoretical development phase with increasing depth and structure to knowledge representation and inferencing schema. Inductive inference systems have entered the empirical phase where there are beginning to be available a range of usable tools. Autonomous activity systems are still in the research phase and the lack of experience and foundations will restrict the development of knowledge acquisition systems for goal-directed, planning activities during this generation.

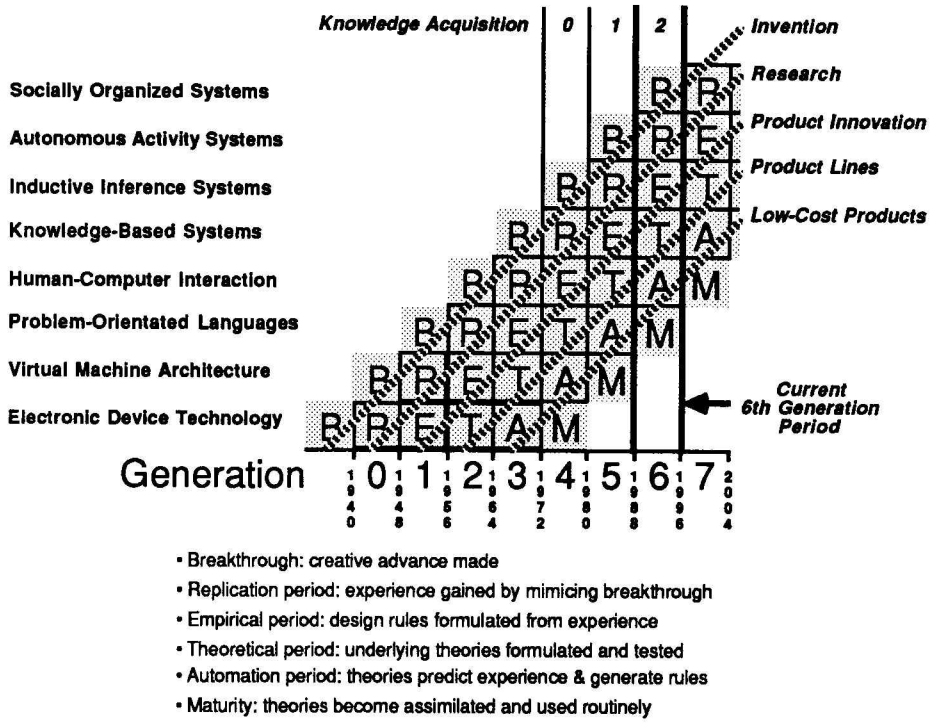


Fig.3 Second Generation Knowledge Acquisition Systems in the Infrastructure of Information Technology

Thus overall, we should expect the emphasis in second generation knowledge acquisition systems to be on the use of high-speed graphic workstations, programmed in high-level languages, using high-quality user interfaces to provide effective knowledge representation, some inductive inference and little or no capability to deal with planning.

ELICITATION FROM MULTIPLE EXPERTS

The knowledge to be acquired for the development of knowledge based systems is often distributed across a group of experts rather than available for elicitation from a single expert. Group elicitation presents major problems because experts can disagree on the use of concepts and vocabulary, and this disagreement may be tacit causing confusion. Methodologies and tools are being developed for knowledge acquisition from groups in which the conceptual frameworks of different experts are compared in a way that makes such disagreements overt and readily identifiable (Shaw & Woodward 1988, Shaw 1988, Gaines & Shaw 1989). Studies have also been made of extending performance systems based on group elicitation to allow multiple knowledge structures from different experts to be combined without distortion, enabling "dissenting opinions" to be requested (Boose 1986).

Apparent dissent between experts may be due to conflict in using the same terminology for different concepts. In addition apparent difference in viewpoints can result from the use of different terminology for the same concept. These sources of dissent between experts apply equally to the conceptual frameworks of experts and clients. Figures 4 and 5 below illustrate some of the phenomena that can arise by considering the mappings from the concepts of experts or clients to the relevant *distinction* in the domain and the *term* used for that distinction.

In Figure 4: *consensus* arises if the conceptual systems of the two individuals assign the same term to the same distinction; *conflict* arises if they assign the same term to different distinctions. In Figure 5: *correspondence* arises if the conceptual systems of the two individuals assign different terms to the same distinction; *contrast* arises as they assign different terms to different distinctions. *Conflict* shows possible dissent over the application of a concept in a domain since the same term is being used in different ways. *Correspondence* shows possible dissent over the appropriate term to be used for a concept. *Consensus* is particularly important in showing a shared conceptual sub-system. *Contrast* shows conceptual sub-systems that are not shared but do not lead to dissent.

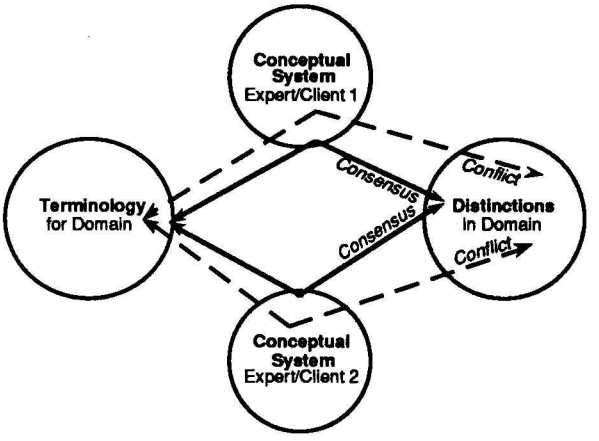


Fig.4 Consensus and conflict between conceptual systems

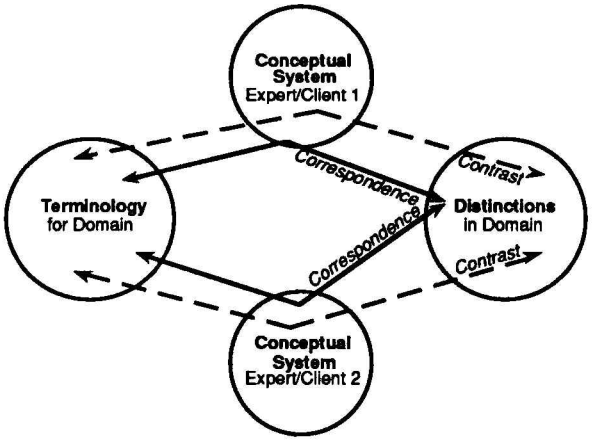


Fig.5 Correspondence and contrast between conceptual systems

IMPROVED VALIDATION

The validation of knowledge-based systems is a complex topic in its own right, and a variety of techniques and tools are being developed to support different forms of validation (Shaw & Woodward 1988, Benbasat & Dhaliwal 1988). The objective validation of the knowledge base in terms of its performance, its subjective validation in comparison with experts, and its referential validation in comparison with other knowledge sources should all be supported. Objective validation may require significant case histories that can be collected only from an operational system in actual applications. Referential validation may require deep analysis of derivations of the knowledge base from other knowledge bases. There is a natural separation in subjective validation between experts' direct agreement that their rationale corresponds to that of the knowledge base,

and observed correspondence between their behavior and that of the system based on the knowledge base. Both types of validation should be supported.

The validation of knowledge-based systems has all the dimensions of normal scientific validation together with those concerned specifically with the emulation of expertise (Gaines 1988b). As shown in Figure 6, the system can be validated through four primary pathways:

- Through its correspondence to actual events checked against empirical data;
- Through its efficacy in applications checked against actual applications;
- Through the adequacy of its foundations checked against other disciplines corresponding to the derivation of presuppositions from other disciplines—this is sometimes termed 'deep knowledge';
- Through the comprehensibility of explanations given checked against the relevant reference community—corresponding to explanation.

The expert systems paradigm introduces the two additional pathways shown on the right of Figure 6 by taking into account the possibility that there may be "experts" who are competent in applications and explanations:

- Through its correspondence to expert performance checked against the behavior of experts in applications
- Through its correspondence to expert knowledge checked against the rationale of experts in communicating with their reference community

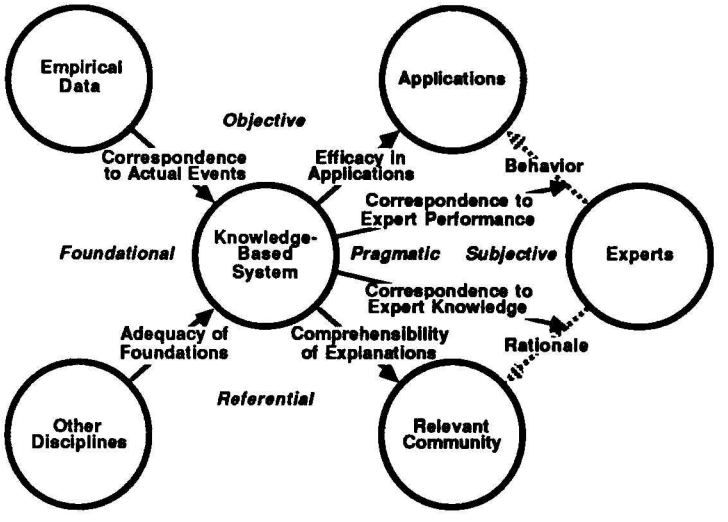


Fig.6 Dimensions of validation of knowledge-based systems

INTEGRATION OF MACHINE LEARNING

There is currently a major paradigm split in knowledge acquisition research and practice between techniques for the transfer of existing knowledge from human experts and those for the creation of new expertise through machine learning. There is, however, a fundamental relation between the two paradigms in that existing expertise was at some time derived through empirical induction. There is also continuity between the two paradigms in that existing expertise may be partial, erroneous, and of various forms, such that it cannot completely replace empirical induction but may serve to to guide and expedite it. For example human experts in a decision task might know: at one extreme, a complete, minimal set of correct decision rules; through a critical set of cases described in terms of a minimal set of relevant attributes with correct decisions; to, at the other extreme, a source of cases described in terms of attributes among which are those relevant to the decision and with only a greater than chance probability of correct decisions. This is a sequence of decreasing knowledge on the part of the human expert.

Gaines (1989) has studied how the amount of knowledge available from the expert affects the amount of data required for effective empirical induction in the range of situations above. He took a set of cases originally defined by Cendrowska (1987) as a test of empirical induction and generated a range of datasets from it by making incorrect decisions with known probability and adding a prescribed number of irrelevant attributes with random values. He then applied a powerful empirical induction algorithm, INDUCT, to the datasets to determine the average amount of data required to obtain a correctly performing set of rules. The original dataset can be correctly modeled with 6 rules involving 4 attributes.

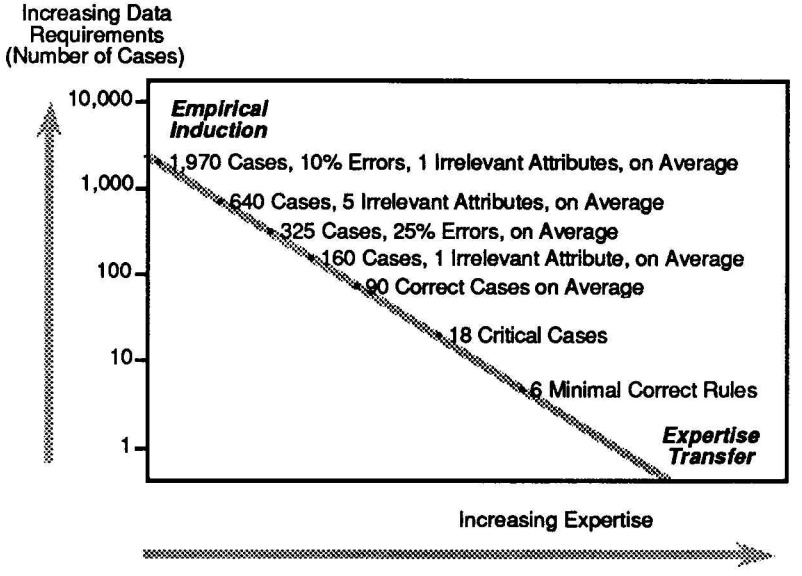


Fig.7 The Trade Off Between Data and Knowledge

Figure 7 shows some of Gaines' results as a plot of the data required on average for empirical induction against the type of knowledge available from the simulated expert. The trade off between data and expertise is clearly shown, and what is also apparent is the continuum between dependence on empirical induction from cases on the left and dependence on transfer of expert knowledge on the right. It is particularly noteworthy that a single algorithm was used for empirical induction across the range of situations Gaines investigated. The results validate the claims of those with tools aimed primarily at eliciting relevant attributes and critical cases from experts that this is an efficient approach. They also show that empirical induction and expertise transfer can be combined to produce a more powerful and versatile acquisition tool than either alone. In future we can expect methodologies and tools to increasingly support a combination of techniques.

INTEGRATION OF TEXT ANALYSIS

A very significant component of human knowledge has already been expressed in text, diagrams and pictures and captured in books and journals. The extraction of knowledge from such material is a major research area in its own right which has major relevance to knowledge acquisition and may be expected to produce methodologies and tools that will be integrated with those described in this tutorial.

One class of research in concerned with the detailed analysis of knowledge structures in text. The CYC project at MCC is a notable example of an attempt to build a knowledge base from an encyclopaedia (Lenat, Prakash & Shepherd 1986). Slator (1988) has described a system for producing text-specific lexicons from selected machine-readable dictionary definitions.

Szpakowicz (1988) has described a system that semi-automatically processes technical text and incrementally builds a conceptual model of a domain using Sowa's (1984) conceptual graphs. More simplistic forms of text analysis have been incorporated in a number of knowledge elicitation tools as means of pump-priming the elicitation with a range of entities and concepts. KSS0 (Gaines & Shaw 1987), KITTEN (Shaw & Gaines 1987), KRITON (Linster 1988) and Cognosys (Woodward 1988) all provide tools for text presentation and semi-automatic analysis as part of the knowledge elicitation process.

As techniques for knowledge extraction from text are developed they will play an increasingly important role in knowledge acquisition. It is unlikely that textual sources will ever replace elicitation from experts for many of the reasons already discussed—knowledge is continually being created and experts, if they write it down at all, do so after considerable delay and with much loss of significant components. However, incorporation of what knowledge is available from written media will become an important source. The same considerations apply to diagrams where, however, the state of the art of knowledge extraction is rather less advanced than with text.

CONVERGENCE WITH HYPERMEDIA

Hypertext and hypermedia are computer-based extensions of existing media such as books and films, and the representation of knowledge in these media has important applications to knowledge acquisition (Nelson 1980, Conklin 1986). Hypermedia are particularly important in knowledge acquisition because much human knowledge is implicit in action and not formally represented and processed. Hypermedia go beyond knowledge bases in not requiring that the information system be able to represent and process the knowledge encoded in the media used. They subsume knowledge bases but extend knowledge representation to include any media whereby human knowledge can be transferred even if there is no means for analyzing that knowledge or its transfer. For example, a film of a skilled person performing a task which enables another person to perform that task better contains knowledge which can be accessed and transferred through hypermedia (Vickers & Kingston 1987).

A number of knowledge acquisition tools have already been written in Apple's HyperCard and this trend may be expected to continue. Developers of knowledge-based systems shells are providing links from them to hypermedia products to support the application of the wide variety of knowledge representations offered. Knowledge acquisition tools are being integrated with hypermedia systems for similar reasons, and this raises interesting possibilities for the representation and use in explanation and tutoring of a much wider range of forms of knowledge than has been possible in the past. The support of knowledge acquisition methodologies through hypermedia tools and the integration of knowledge acquisition systems, hypermedia systems and knowledge-based systems shells is a major industry trend.

CONVERGENCE WITH CASE

Computer-Aided Software Engineering (CASE) tools range from those that are clearly technical and logical extensions of existing compiler and support environment techniques to those which attempt to deal with many of the "softer" issues of system development such as tracing customer requirements and design decisions. There are many levels of system development for which there now exist structured development methodologies and CASE tools and these are equally applicable to the development of knowledge-based systems. It is interesting to examine potential applications of CASE tools in knowledge-based system development, and the relation between knowledge acquisition tools and similar CASE tools. For example, data design tools using entity-attribute-relation models are related to the entity-attribute, or repertory, grids used in many knowledge elicitation tools.

Figure 8, as a system development model, makes one realize how many aspects are neglected in current knowledge-based system developments. AI technologies are as yet highly immature compared with conventional software engineering. In particular, knowledge base maintenance and system failure analysis will become major problem areas as significant systems come into use. However, knowledge acquisition tools designed to work directly with experts and support them through non-technical interfaces show up very well compared with CASE tools eliciting the same knowledge. The two families of tools have been designed for different markets and user communities.

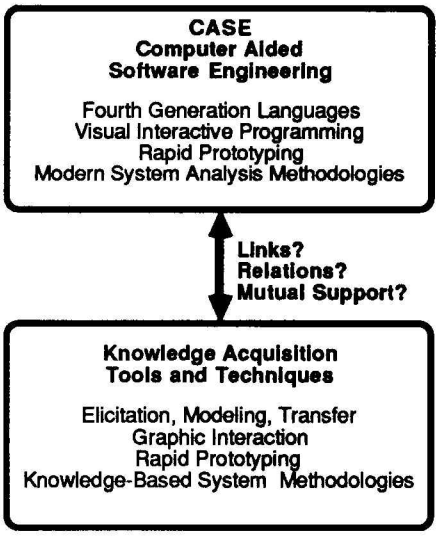


Fig.8 Synergisms between CASE and Knowledge Acquisition

CONCLUSIONS

There are many possible approaches to knowledge engineering. However, most researcher and practitioners agree that knowlege acquisition is the bottleneck in the task of developing knowledge based systems. It has been shown in this paper how the role of the knowledge engineer is changing from the repository of expertise to the manager of the whole elicitation and encoding process. Several automated methods already exist, and in this paper some well-established methods of eliciting knowledge used in other disciplines are shown to be relevant to the knowledge acquisition problem. The underlying technological infrastructure was put into the perspective of computer science developments, and related to the social phenonemon of expertise. Some future developments in integrating knowledge acquisition with existing methodologies in other areas were elucidated such as dealing with the problems caused by having to use multiple experts; the problems of validation of knowledge acquisition and knowledge based systems; the use of tools from the field of machine learning; the relationship with hypermedia systems; and the convergence with CASE methodologies.

ACKNOWLEDGEMENTS

Financial assistance for this work has been made available by the Natural Sciences and Engineering Research Council of Canada.

REFERENCES

Benbasat, I. & Dhaliwal, J.S. (1988) A framework for the validation of knowledge acquisition. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Third AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.1-1-1-18. Banff .

Boose, J. H. (1986) Rapid Acquisition and Combination of Knowledge from Multiple Experts in the Same Domain. **Future Computing Systems 1**(2) 191-216

Cendrowska, J. (1987) An algorithm for inducing modular rules. **International J. of Man-Machine Studies 27** (4), 349-370 .

Conklin, J. (1986). A survey of hypertext. MCC Technical Report STP-356-86.Austin:Microelectronics & Computer Technology Corporation.

Gaines, B.R. (1988a). Second generation knowledge acquisition systems. **Proceedings of the Second European Workshop on Knowledge Acquisition for Knowledge-Based Systems (EKAW'88)**. GMD-Studien Nr.143. pp.17-1-17-14. Bonn, Germany: Gellsellschaft für Mathematik und Datenverarbeitung (June).

Gaines, B.R. (1988b) Knowledge acquisition and technology. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Third AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.8-1-8-21. Banff (November).

Gaines, B.R. (1989) An Ounce of Knowledge is Worth a Ton of Data: Quantitative Studies of the Trade-Off between Expertise and Data based on Statistically Well-Founded Empirical Induction. KSI-89-2, Knowledge Science Institute, University of Calgary.

Gaines, B.R. & Shaw, M.L.G. (1986). A learning model for forecasting the future of information technology. **Future Computing Systems, 1**(1), 31-69.

Gaines, B.R. & Shaw, M.L.G. (1987). Knowledge support systems. **ACM MCC-University Research Symposium**. Austin, Texas: MCC. pp.47-66.

Gaines, B.R. & Shaw, M.L.G. (1989) Comparing the conceptual systems of experts. **Proceedings of the Eleventh International Joint Conference on Artificial Intelligence**, Detroit, Michigan, USA, 633-638.

Lenat, D., Prakash, M. & Shepherd, M. (1986). CYC: Using common sense knowledge to overcome brittleness and knowledge acquisition bottlenecks. **AI Magazine 6**(4), 65-85.

Linster, M. (1988) A critical look at KRITON. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Third AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.17-1-17-19. Banff (November).

Nelson, T. (1980). Interactive systems and the design of virtuality. **Creative Computing, 6**(11), 56-62 (November).

Shaw, M.L.G. (1988). A framework for knowledge acquisition from more than one experts. **Proceedings of the Third IASTED International Conference,** Los Angeles, 19-22.

Shaw, M.L.G. & Gaines, B.R. (1987) KITTEN: Knowledge Initiation & Transfer Tools for Experts & Novices. **International Journal of Man-Machine Studies 27**(3), 251-280 (September).

Shaw, M.L.G. & Woodward, J.B. (1988). Validation in a knowledge support system: construing and consistency with multiple experts. **International J. of Man-Machine Studies, 29,** 329-350.

Slator, B.M. (1988) Extracting lexical knowledge from dictionary text. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Third AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.33-1-33-18. Banff (November).

Sowa, J. (1984) **Conceptual Structures**. Addison Wesley Publishing Co. Reading Ma. 1984.

Steels, L. (1986). Second generation expert systems. Bramer, M.A. (Ed). **Research and Development in Expert Systems III**. pp.175-183. Cambridge: University Press.

Szpakowicz, S. (1988) Semi-automatic acquisition of conceptual structure from technical texts. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Third AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.35-1-35-16. Banff.

Vickers, J.N. & Kingston, G.E. (1987). Modelling the master coach: building an expert system for coaching. Norrie, D., Ed. **Learning in Future Education: Proceedings of International Conference on Computer Assisted Learning in Post-Secondary Education**. pp.207-212. University of Calgary.

Woodward, B. (1988) Knowledge engineering at the front-end: defining the domain. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Third AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.37-1-37-25. Banff (November).