ELICITING AND COMPARING THE CONCEPTUAL STRUCTURES OF EXPERTS AND END-USERS

Mildred L G Shaw & Brian R Gaines
Knowledge Science Institute
Department of Computer Science
University of Calgary
Calgary, Alberta, Canada T2N 1N4
mildred@cpsc.ucalgary.ca, gaines@cpsc.ucalgary.ca

Abstract: Knowledge acquisition tools developed for the rapid prototyping of expert systems are, at one level, CASE tools for knowledge-based systems and, at another level, requirements specification elicitation tools applicable to a wide range of system developments. This paper briefly reviews trends in knowledge acquisition tools and methodologies since our presentation at CASE'88, and then presents, in depth, knowledge acquisition methodologies and tools concerned with eliciting and comparing the conceptual structures of experts and end-users. Such tools have general applications in CASE, for example, in obtaining and refining initial requirements specifications.

INTRODUCTION

At CASE'88 in Boston we reported on trends in software engineering for knowledge-based systems (Gaines 1988b), and noted that the methodologies and tools being developed by the computer-based "knowledge acquisition" community were essentially CASE tools designed for use by end-users. In the paper and presentation we demonstrated this with examples from KSS0 (Gaines & Shaw 1987), an integrated suite of knowledge acquisition tools for interviewing experts, analyzing text, and inductively modeling case histories. At CASE'89 a complete session is being devoted to knowledge acquisition with further examples of state-of-the-art tools such as Aquinas (Boose & Bradshaw 1987) from Boeing Computer Services.

The technical section of this paper focuses on the tools in KSS0 that enable the conceptual structures of experts and clients to be compared and contrasted, since these tools seem to have general application in CASE, for example, in obtaining and refining initial requirements specifications. The initial section focuses on discussions within the knowledge acquisition community about trends in the methodologies and tools being developed (Shaw & Gaines 1989).

TRENDS IN KNOWLEDGE ACQUISTION TECHNOLOGIES

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 (Boose & Gaines 1988, Gaines & Boose 1988). 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 (Gaines 1987a,b,c, 1988b). 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.

Steels (1986) has suggested that we are entering an era of *second generation* expert systems, "able to combine heuristic reasoning based on rules, with deep reasoning based on a model of the problem domain. This solves a number of major problems of current expert systems, most importantly the problem of knowledge acquisition: second generation expert systems can learn new rules by examining the results of deep reasoning." If one generalizes this concept and thinks of second generation systems as having *embedded knowledge acquisition* using a variety of sources including those suggested by Steels, then it captures the thrust of current research: to integrate knowledge acquisition and performance systems. The zeroth generation of manual knowledge acquisition techniques led to the first generation of computer-aided tools, and this is now leading to a second generation of integrated knowledge support systems.

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

Figure 1 identifies the major trends in knowledge acquisition research and tool development leading to the development of second generation systems. The remainder of this paper focuses on the third topic down, the elicitation and comparison of conceptual structures from multiple experts.

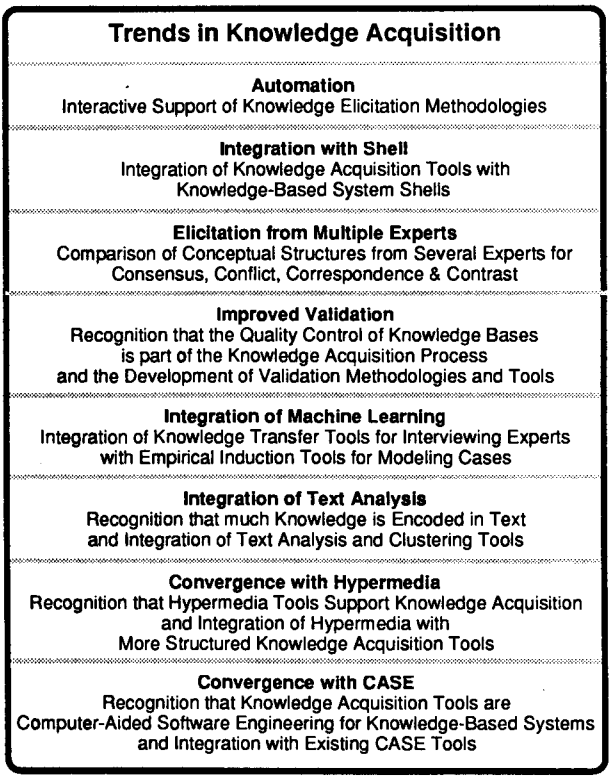

Fig.1 Trends in knowledge acquisition research and tool development

SOURCES OF DISSENT IN CONCEPTUAL SYSTEMS

Computer elicitation of entity-attribute, or *repertory*, grids has proved to be a powerful technique for acquiring the declarative knowledge structures of an expert in a domain (Shaw & Gaines 1983, 1987, Boose 1984, Boose & Bradshaw 1987, Diederich, Ruhmann & May 1987, Gaines 1987c). Studies have also been made of extending performance systems based on these grids to allow multiple knowledge structures from different experts to be combined without distortion, enabling "dissenting opinions" to be requested (Boose 1986). In some cases, however, apparent dissent between experts is 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 2 and 3 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 2: *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 3: *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.

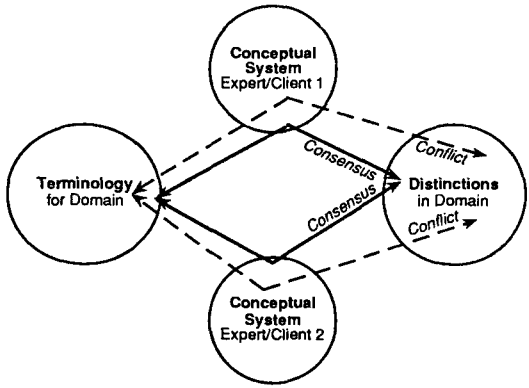

Fig.2 Consensus and conflict between conceptual systems

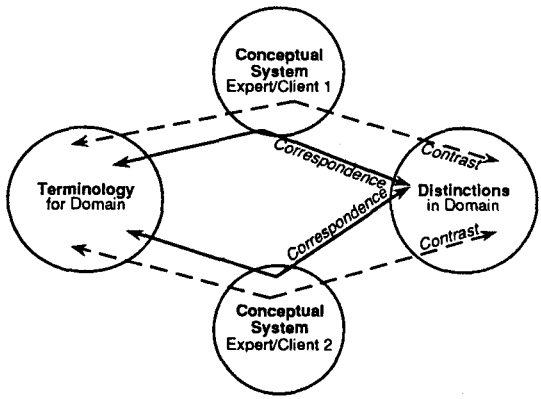

Fig.3 Correspondence and contrast between conceptual systems

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. The following section describes a methodology for eliciting and processing entity-attribute grids that derives this classification for groups of experts and clients.

COMPARING CONCEPTUAL SYSTEMS

Any comparison of conceptual systems necessarily involves approximation since a complete conceptual system may involve indefinitely complex relations and different concepts will never be identical in all respects. However, in the initial phases of knowledge acquisition, highlighting gross similarities and differences is itself valuable in promoting directed discussion among experts and clients that can lead to the exposure of more subtle relationships. As a start one wishes to elicit the major distinctions that an individual uses in a domain, the terminology used for them, and the relation of such distinctions and terminology to those of others.

Entity-attribute grid elicitation is an effective method for eliciting major distinctions and terminology in a domain (Shaw 1980). It is an extensional approach in which individuals are asked to specify a set of entities in a domain, then make distinctions among them, naming the distinctions and classifying all the specified entities in terms of them. The extension of a distinction determined in this way is only an approximation to the underlying concept since critical entities may be missing in the classification. However, both manual and computer-based elicitation techniques attempt to prompt the individual for additional entities to discriminate between extensionally related distinctions (that is making the same, or similar, classifications). Group comparisons, as discussed in this paper, have similar dynamics—an extensionally apparent consensus or correspondence may be accepted or rejected, and the rejection may be based on the specification of additional entities as counter-examples. Knowledge acquisition is essentially a negotiation process leading to approximations to conceptual structures that are adequate for some practical purpose such as system development.

Measures for comparing extensionally defined distinctions may be formulated in terms of a fuzzy sets model of grid data (Gaines & Shaw 1980, 1986). To apply a measure of similarity between distinctions it is necessary that they have the same extension—that is, that the group agrees on a common set of entities among which they will each make distinctions. Phase 1 of the methodology consists of establishing this common set of entities and ensuring, as far as possible, that each individual in the group is defining the entities in the same way. This appears straightforward if the entities are well-defined, named concrete objects in the world. However, the concept of an entity turns out to have its own difficulties in practice—if the entities are people, is it important in what mood and what role they are?—if the entities are tools, is it important who is using them and to what purpose? Again, entity definitions are part of the ongoing negotiation process of knowledge acquisition, and tacit disagreement over them can be a source of dissent. The methodology described compares entities also using a similar measure over the attributes being used.

For the sake of simplicity, we will phrase the following in terms of comparisons between pairs of individuals. The methodology extends simply to a group as discussed later. Correspondence and contrast between two conceptual systems may be measured by eliciting independent grids from the two individuals and using the similarity measure to determine for each distinction in one grid the best matching distinction in the other. Distinctions which match above an upper threshold are in correspondence and those that match below a lower threshold are in contrast—those matching between the thresholds are ambiguously related.

The measurement of consensus and conflict requires a common terminology for the distinctions and may be based on the *exchange* of grids. One individual is given the grid elicited from the other with only the attributes and entities specified, not the values, and asked to rate the entities on the attributes. The similarity measure can then be used to determine how well the distinctions made by the second individual match with those made by the first. Distinctions which match above an upper threshold are in consensus and those that match below a lower threshold are in conflict—those matching between the thresholds are ambiguously related.

As important as the methodology is the presentation of the results to the individuals in such a way that they readily visualize the areas of potential dissent and their sources. Figure 4 is a graphic presentation of a consensus and conflict analysis from a grid exchanged between two experts in geographic mapping techniques (Shaw & Woodward 1987). The differences are shown in the grid on the left in numeric form and as shading with darker gray indicating greater difference. The sets of attributes and entities have both been sorted so that those with highest matches indicating greatest consensus are at the top of their respective lists. The graphs on the right show the declining match values and may be used to decide where to place thresholds distinguishing consensus and conflict. Most

importantly, this presentation contains sufficient raw and processed data that it may be used for a range of discussions, from simple identification of apparent disagreements, such as the experts disagree on the use of the term "linear interpolation" and the entity "probability matching," to a detailed analysis of how these disagreements have arisen. Measures of correspondence and contrast are presented in a similar graphical form with attributes paired from the two grids.

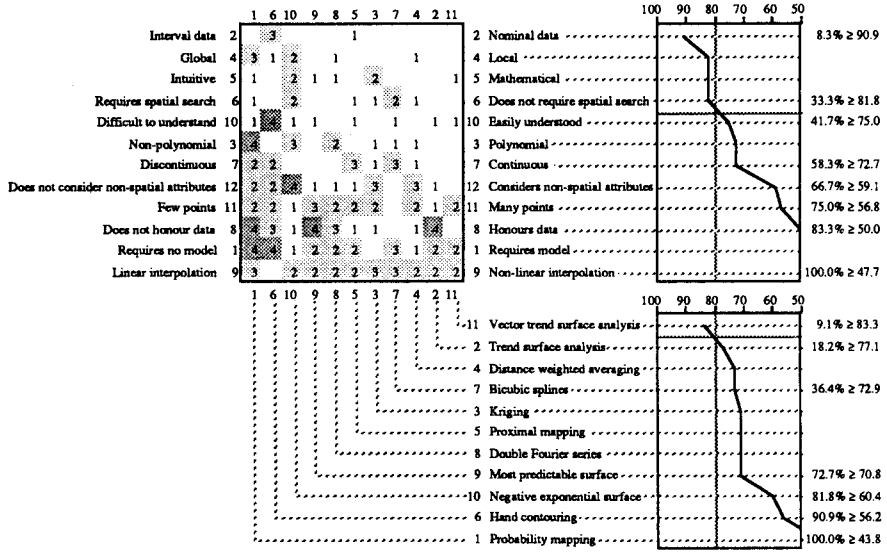

Fig.4 Presentation of measured consensus and conflict

The results of the comparisons are summarized and presented in the four quadrant diagram exemplified in Figure 5. The *consensual* items are simple to deal with—an obvious question is "to what extent can a knowledge-based system be built based only on these well-agreed distinctions?" The *conflicting* items can be looked at with the reverse implication, "to what extent can a knowledge-based system be built not using these ill-defined distinctions?" If some of the distinctions are agreed to be necessary for system development then the basis of the conflict can be examined in greater detail starting with the presentation of Figure 4. Many resolutions of the conflict are possible, but the computer-based methodology can go no further than making the basis of conflict and the possible need for resolution apparent.

The *corresponding* items may again be looked at in terms of their relevance to the problem. In addition their further interpretation often leads to significant insights into the knowledge structures involved. For example, the terms *low-level-data—high-level data* and *nominal data—interval or ratio data* corresponding may be interpreted as arising from different *levels of abstraction*. Questions then arise about whether the levels should be represented in the knowledge structures, "which will be best understood by a client?", and "should one term be used as an explanation of the other in response to a *help* request?" Interpretation of the correspondences in Figure 5 is an important knowledge acquisition methodology in its own right. The simple example shown brings to light abstraction hierarchies, physical and social causal links, and other relations between attributes. The methodology has the advantages of group *brainstorming* in stimulating knowledge elicitation, and captures the results in a highly structured and usable form.

Note that the derivation of consensual, corresponding, conflicting and contrasting attributes is completely algorithmic, based solely on the data in the grids. This derivation is done by a computer program, not a knowledge engineer, and its basis can be demonstrated clearly to the experts and clients through computer output such as Figure 4. Thus, there are no opinions being expressed about the correctness of the use of the attributes and terminology,

that the differences highlighted are 'right' or 'wrong.' It is open to the experts and clients to consider, discuss and explain these differences, changing or retaining them as they wish. Conflicts can be retained in the final system if desired by tagging classes, objects and rules with the sources from which they derive.

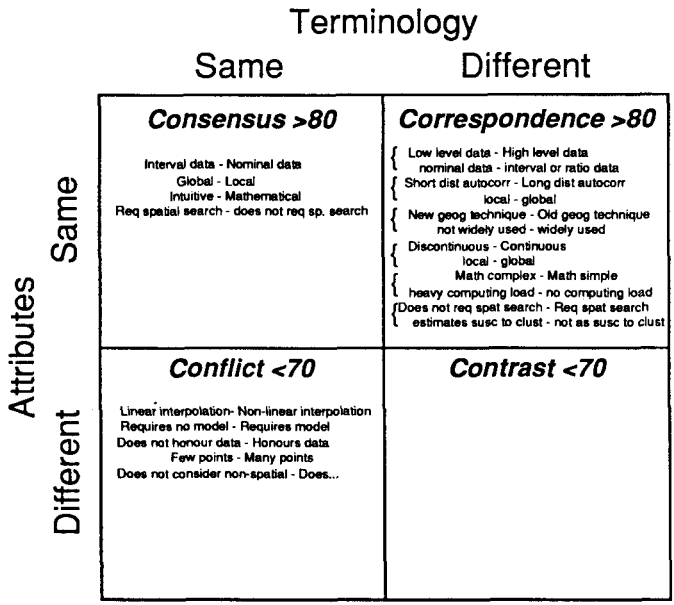

Fig.5 Summary of comparisons between two experts

The methodology has been described in relation to comparisons between two people. However, it scales up linearly with the size of the group, so that n people will be involved in n elicitations, one base elicitation and n-1 exchanges. The following section reiterates the methodology in detail as a practical technique for knowledge acquisition with groups of experts and clients.

A METHODOLOGY FOR COMPARING CONCEPTUAL SYSTEMS

Figure 6 shows the phases of a methodology for comparing conceptual systems:

Phase 1: Domain Discussion and Instantiation

- The problem is discussed with the group of experts and representative clients, and a domain is identified.
- The use of computer-based knowledge acquisition is introduced, a specific task is agreed upon, and a purpose for the elicitation developed.
- The first step in acquisition consists of the elicitation of a grid from each individual. This gives each person experience in using the computer-based system and provides an initial set of entities and attributes particular to the individual.
- These grids are processed using graphical clustering methods, and the results shown and explained to the individual from whom the data was elicited. As a result an individual may extend or elaborate their conceptual system expressed in the grid.
- The group meets together to discuss each set of entities, review and revise the purpose or context of the elicitations, and agree on a common purpose and set of entities which all understand to be used in the next two phases.

Phase 1: Domain Instantiation

Experts together agree a set of entities relevant to the domain

Phase 2: Conceptualization

Experts individually conceptualize entities in terms of attributes and values

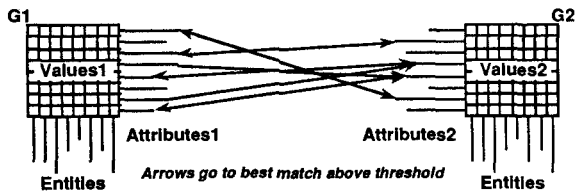

Attributes with arrows between show **Correspondence**
Attributes with no arrows between show **Contrast**

Phase 3: Exchange

Experts individually conceptualize entities in terms of each other's attributes

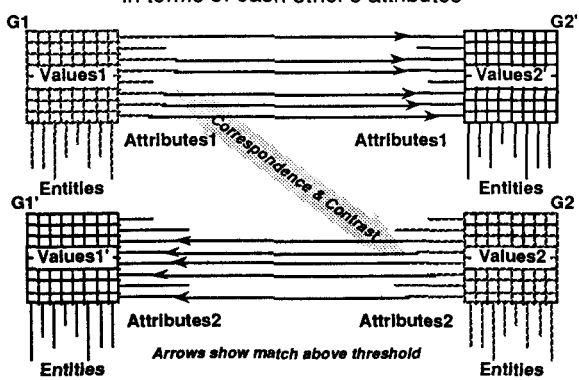

Attributes with arrows between show **Consensus**
Attributes with no arrows between show **Conflict**

Fig.6 A methodology for comparing conceptual systems

Phase 2: Conceptualization and Feedback

- This common entity set provides the basis for the next phase of knowledge acquisition which consists of the elicitation of a grid from each individual using the set of common entities.
- These grids are again graphically clustered, and the results reported and explained to each expert respectively. Again this may lead to extension and elaboration.
- This data is then processed to identify corresponding and contrasting attributes.

Phase 3: Exchange and Compare

- The next phase of knowledge acquisition consists of the individuals exchanging their grids elicited in phase 2 with each other individual in the group to produce grids in which one person classifies the common set of entities using another person's attributes.
- This data is then processed to identify consensual and conflicting attributes.
- Their conceptual systems and the relations between them are then discussed by the group. Consensus is identified. Conflict is resolved, qualified or noted as irresolvable. Correspondence is interpreted as due to differences in terminology or correlations through some identified phenomenon. Contrasts are analyzed as irrelevant or a relevant part of some specialist viewpoint.
- This analysis is propagated through all further phases of knowledge acquisition such as class, property and object definitions, and rule specification. For example, possible conflicts are avoided if possible, and identified and explained otherwise—correspondences are used to generate alternative forms of questions and explanations.

CONCLUSIONS

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. This paper describes a computer-supported methodology 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.

The methodology provides facilities for revealing the similarities and differences in the conceptual systems of groups of experts and clients. It can be used to focus discussion between experts, and between them and representative clients, on those differences between them which require resolution, enabling them to classify them in terms of differing terminologies, levels of abstraction, disagreements, and so on.

The methodology promotes the full exploration of the conceptual framework of a domain of expertise by encouraging experts and clients to operate in a *brainstorming* mode as a group, using differing viewpoints to develop a rich framework. It avoids social pressures forcing an invalid consensus by providing objective analysis of separately elicited conceptual systems.

ACKNOWLEDGEMENTS

Financial assistance for this work has been made available by the Natural Sciences and Engineering Research Council of Canada. We are grateful to many colleagues at the knowledge acquisition workshops for discussions which have highlighted many of the issues raised in this paper. We are particularly grateful to Brian Woodward, John Boose and Jeff Bradshaw for many fruitful discussions of knowledge acquisition systems.

REFERENCES

Boose, J.H. (1984). Personal construct theory and the transfer of human expertise. **Proceedings AAAI-84**, 27-33. California: American Association for Artificial Intelligence.

Boose, J.H. (1986). Rapid acquisition and combination of knowledge from multiple experts in the same domain. **Future Computing Systems**, 191-216.

Boose, J.H. & Bradshaw, J.M. (1987) Expertise transfer and complex problems: using AQUINAS as a knowledge acquisition workbench for knowledge-based systems. **International Journal of Man-Machine Studies(26)**, 3-28.

Boose, J.H. & Gaines, B.R., Eds. (1988). **Knowledge Acquisition Tools for Expert Systems**. London, Academic Press.

Diederich, J., Ruhmann, I. & May, M. (1987) KRITON: a knowledge-acquisition tool for expert systems. **International Journal of Man-Machine Studies(26)**, 29-40.

Gaines, B.R. (1987a) An overview of knowledge acquisition and transfer. **International Journal of Man-Machine Studies** 26(4), 453-472 (April).

Gaines, B.R. (1987b) Advanced expert system support environments. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Second AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.8-0-8-14. Banff (October).

Gaines, B.R. (1987c). Rapid prototyping for expert systems. Oliff, M.D., Ed. **Intelligent Manufacturing: Proceedings from First International Conference on Expert Systems and the Leading Edge in Production Planning and Control**. pp.45-73. Menlo Park, California, Benjamin Cummins.

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). Software engineering for knowledge-based systems. **Proceedings of CASE'88, Second International Workshop on Computer-Aided Software Engineering**. 14-1-14-7.

Gaines, B.R. (1988c) 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. & Boose, J.H., Eds. (1988). **Knowledge Acquisition for Knowledge-Based Systems**. London, Academic Press.

Gaines, B.R. & Shaw, M.L.G. (1980). New directions in the analysis and interactive elicitation of personal construct systems. **International Journal of Man-Machine Studies** 13(1) 81-116 (July).

Gaines, B.R. & Shaw, M.L.G. (1986). Induction of inference rules for expert systems. **Fuzzy Sets and Systems**, 8(3), 315-328 (April).

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

Shaw, M.L.G. (1980). **On Becoming a Personal Scientist**. London: Academic Press.

Shaw, M.L.G. & Gaines, B.R. (1983). A computer aid to knowledge engineering. **Proceedings of British Computer Society Conference on Expert Systems**, 263-271 (December). Cambridge.

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. & Gaines, B.R. (1989). Knowledge acquisition: some foundations, manual methods and future trends. **Proceedings of the Third European Workshop on Knowledge Acquisition for Knowledge-Based Systems (EKAW'89)**. Paris, France (June).

Shaw, M.L.G. & Woodward, J.B. (1987) Validation of a knowledge support system. Boose, J.H. & Gaines, B.R. (Eds) **Proceedings of the Second AAAI Knowledge Acquisition for Knowledge-Based Systems Workshop**. pp.18-0-18-15. Banff (October).

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