University of Groningen

How to improve our understanding of group decision making with the help of artificial intelligence

Liebrand, Wim B.G.

Published in:
Acta psychologica. International journal of psychonomics

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
1992

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Liebrand, W. B. G. (1992). How to improve our understanding of group decision making with the help of artificial intelligence. *Acta psychologica. International journal of psychonomics*, 80(1-3), 279-295.

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

REPRINTED FROM:

International Journal of Psychonomics

acta psychologica

Volume 80, Nos. 1-3, 1992

How to improve our understanding
of group decision making with the help
of Artificial Intelligence

Wim B.G. Liebrand *

pp. 279-295

North-Holland — Amsterdam

How to improve our understanding of group decision making with the help of Artificial Intelligence

Wim B.G. Liebrand *

University of Groningen, Groningen, The Netherlands

Within science we primarily obtain knowledge of a specific field by reading the published results of theoretical and empirical studies. It is argued that this approach may lead to a biased and incomplete perspective of a research area. It is proposed to also use methods from Artificial Intelligence to elicit, model and use the knowledge of experts. It is expected that especially their heuristic knowledge is relevant for problem solving and consultancy applications. An illustrative example concludes the paper.

Understanding group decision making has proven to be an important and complex task for many researchers within various disciplines. Previous research has described both the ways in which interdependent people go about solving problems and making decisions as well as how they should have made them.

In the area of group decision making we observe major departures from what in terms of Subjective Expected Utility (SEU) theory would be called rational decision making. The elegant mathematical world of decision making obviously does not sufficiently match reality. In light of the rather stringent assumptions being made it seems not very surprising that we cannot call man an 'SEU-rational' decision maker. People normally do not have available complete information about all utility functions involved, and as a consequence they are not in the position to carry out the required numerical operations. What is more,

* I am grateful to D.M. Messick, T. Takigawa and E. Boekkooi-Timminga for their support in completing this paper.

Correspondence to: W.B.G. Liebrand, Social Science Information Technology, University of Groningen, Grote Rozenstraat 15, 9712 TG Groningen, The Netherlands.

even if all the necessary information were available, their limited capacity for information processing would impede such operations.

On the other hand, empirical research on group decision making and problem solving shows that people transform the structure of the decision setting according to their individual outcome preferences, their short- and long-term strategic goals, and their selective perception of complex structures. These numerous and diverse empirical research findings, however, add to the complexity of the field and as a consequence we currently lack a comprehensive theory of decision making in interdependent settings, which integrates the variety of theoretical and empirical research findings. This state of affairs, of course, is undesirable both for the understanding of decision making and in a more applied sense for consultancy applications.

The major aim of the present paper is to provide a possible way to integrate previous research findings in a pragmatic way, thereby allowing researchers and consultants to make better use of those findings. We will try to accomplish that by using recent developments in the field of artificial intelligence. Although the present line of reasoning is applicable to various research areas, we will focus primarily on the domain of decision making in groups in which the well-being of the group members is affected by the decisions taken by the group members. There are two reasons for selecting this domain. First, the author has conducted a few studies in this area. The more important reason, however, is that the study of group decision making relies heavily on game theory and experimental games. The indirect advantage then is that this highly formalized research paradigm could be a perfect domain for the artificial intelligence approach outlined below.

Artificial Intelligence

In the last decade we have observed an ever increasing role of computers in empirical research. Initially, the computer was primarily used to carry out numerical operations, but after the introduction of the personal computer in the early eighties we have witnessed an exponential growth in the capabilities and availability of desktop computers. The computer is now in use in every phase of empirical research and a regular desktop computer is able to carry out millions of instructions per second. This provides a powerful tool for analyzing

highly complex situations. As an example one might think of the allocation of passengers to planes and planes to flight schedules with the objective to minimize costs. The human mind is not very well equipped to perform such a task within a reasonable amount of time.

Despite this 'brute' computational force, and despite several predictions (e.g. Simon and Newell 1958), the computer is still not able to beat the world's chess champion, and it will surely not be able to 'discover and prove an important new mathematical theorem' for a long time. In contrast, experts generally are capable of finding intuitively and almost instantaneously the relevant patterns out of many thousands of different configurations. Experts tend to have a so-called 'tuned awareness' with which they selectively perceive, process, store and recall information. They have strategies for storing and retrieving information or knowledge sets in short- and long-term memory (Tuthill 1990). Access to these strategies would of course be very beneficial to computer applications as well as to non-experts in the field. Next we will discuss this possibility in more detail.

Ordered levels of knowledge

Knowledge may be defined as the result of a process of synthesis in which information is compared to other information and combined into meaningful links (Tuthill 1990). In this view, incoming information is interpreted within an already established personal knowledge framework. This knowledge framework forms a scheme which serves as a model for matching and comparing the incoming information with the current problem set. Within this framework Tuthill distinguishes four levels of knowledge.

Facts are relationships between objects, symbols and events. They can be represented by traditional data structures and are considered to be the lowest level of knowledge. *Concepts* are facts which are grouped with respect to common attributes. Concepts are hierarchically organized in such a way that the lower-level concepts inherit attributes of the ancestor concepts. *Rules* are the next level of knowledge. They operate on facts and concepts. Rules can be conceived of as guides for action and consist of IF-THEN statements representing conditions and actions to take in deductive problem solving. The highest level of knowledge is described in Tuthill's taxonomy as

Heuristic Knowledge. Heuristic knowledge is especially relevant in problem solving. In facing new problem situations, rather than using rule algorithms, shortcuts and combinations of rules are applied. In fact, heuristic knowledge is the personal synthesis of facts, concepts and rules.

The use of heuristic knowledge from experts

The importance of heuristic knowledge seems surprising given the fact that our common understanding of decision making processes consists of facts, concepts and rules. These levels of knowledge are easiest to transfer by means of articles, books and lectures. Heuristic knowledge, or the set of personal rules of thumb, is not normally transferred to others as acquired knowledge. There are several reasons to see this as a highly undesirable situation. The most important of them might be the potential of heuristic strategies to efficiently find solutions to problems. Knowledge about problem-solving strategies is essential because it provides insight into the domain itself. Experts in the field of, for example, group decision making, not only possess knowledge of a large amount of facts, concepts and rules in that research area, they also have heuristic knowledge which transcends rule knowledge. Especially this heuristic knowledge is used in case novel situations are encountered. What is applied in these novel problem situations is a selective synthesis of facts, concepts, and rules previously acquired. All in all they have access to a very complex personal knowledge base which is constantly evolving and which they constantly use in their reasoning about the field. The problem, of course, is that these experts tend to be few in number and mostly not available in case a problem needs to be solved.

The question now arises as to whether it is possible to profit from the heuristic knowledge acquired by experts without having to contact them personally. A tentative answer might be affirmative and related to what 'knowledge engineers' are trying to accomplish. A knowledge engineer tries to emulate expert problem solving in a computer program. In theory their task is simple: elicit knowledge from experts and then organize and represent it in such a way that it can be used by a computer. In practice, knowledge engineers have encountered several problems and 'bottlenecks'.

However, it is only fair to note that substantial progress has been made in this field. As Harmon (1991) points out, knowledge-based systems and expert systems, are slowly recovering from a highly skeptical phase in which, after initial enthusiasm, everybody seemed to know why expert systems could *not* work. Knowledge engineers have now access to several techniques for knowledge elicitation, including structured interviews, protocol analysis, repertory grid and various psychometric scaling techniques. Several methods have also been developed for representing the expert's knowledge in a knowledge-based system. Among the commonly used approaches we find semantic networks, production rules, scripts, frames, repertory grids and object-oriented approaches. All these approaches have their specific advantages and disadvantages. It is not my intention to discuss this in detail (see for example: Gammack 1987; Greenwell 1988; Hart 1989; Parsaye and Chignell 1988). For present purposes it is sufficient to allude to the possibility of representing heuristic knowledge of experts in the field of group decision making.

The domain of group decision making

Decision making in situations of outcome interdependency involves actors whose decisions have reciprocal consequences on their well-being. A convenient way to display different classes of decision making within this domain is shown in fig. 1.

The classification is based on whether the interests of the actors consist of a harmonious set of overlapping interests, a set of competing interests, or a mixture of overlapping and competing interests. The assumed amount of intrapersonal conflict is depicted on the vertical axis. That is, we assume that in case the actors' interests coincide, there will be no ambiguity in what one ought to do, hence actors know for sure that it is in their own interest to cooperate (e.g. sailing a boat). On the opposite side of the horizontal dimension of fig. 1, we again expect low levels of intrapersonal conflict. Here we are dealing with zero sum settings in which gains for one party imply equal losses for the other. Hence, one 'ought' to compete in these situations. Chess playing and the 'battle of the Bismarck Sea' are the often used real-life examples.

Fig. 1. A taxonomy of decision making under settings of outcome interdependency.

Intrapersonal levels of conflict are high in between strictly coordinative and competitive forms of interpersonal relationships. These are situations of partial conflict and involve relationships within which one has the opportunity to both cooperate and compete. Such relationships of partial conflict turn out to be somewhat difficult to describe for reasons of terminology. As Schelling (1973) observes, it is both interesting and puzzling that though we have words to describe individuals in relationships of coordination, e.g., 'partners', 'team members', 'cooperators' and 'collaborators', as well as ones in relationships of strict competition, e.g., 'enemies', 'adversaries', and 'opponents', we do not appear to have similar terms that describe relationships of the partial conflict variety. Strangely, this is the case even though such relationships are highly dominant in wars, strikes, athletic contests, classrooms, relationships between faculty members, and even in marriages. In each of these, the participants find themselves in relationships in which their goals are partly correspondent, and partly in competition with interdependent others. In short, the participants are in a mixed-motive form of relationship. If participants wish to maintain this relationship, they must both receive and deliver rewards to the other, activities generally associated with coordination.

On the other hand, they are also likely to want to obtain the maximum out of a relationship at minimal cost. To the extent that maximum is defined either by the situation, or by their own values as attaining an increased share of a jointly controlled resource at some cost to the other, the relationship also has a strictly competitive or a conflictual component.

This situation in which the decision maker is caught in the dilemma either to cooperate or compete has been a central research topic in the social and behavioral sciences after the important book of Luce and Raiffa (1957). They laid the foundation for the development of experimental games which can be used as research tools to test predictions derived from formal game theory. The most widely used experimental game is the Prisoner's Dilemma Game (PDG). It concerns a situation in which each decision maker is best off acting in his or her own self interest, regardless of what the other persons do. Each self-interested action, however, creates a negative outcome or cost for the other people who are involved. When a large number of people make the self-interested choice, the costs or negative outcomes accumulate creating a situation in which everybody would have done better had they decided not to act in their own short term interest. One of the reasons for its popularity is that the PDG resembles various real life situations of conflict.

An example of a real-life PDG is the decision to pollute. Pollution problems can be found at various levels of decision making, ranging from individuals to nations. For example at the industrial level, no matter what other chemical industries may do to get rid of their chemical waste, it is cheaper to have the waste dumped at some rubbish-dump, or alternatively, in the ocean, than to take care of an adequate solution. The ultimate long-term consequences of this selfish act have to be shared by all individuals. At the individual level the slogan 'every litter bit hurts' nicely reflects the negative consequences accompanying the decision to pollute. Though all individuals would like to avoid the long-term negative consequences, it remains cheaper and simpler for them to keep polluting as anonymous individuals, no matter what the others are doing. As a consequence 'rational behavior' leads to a collective disaster.

The major advantage in using experimental games lies in the amount of control the experimenter has over the decision-making situation. This advantage combined with the possibility to study very

appealing situations of conflict resulted in a veritable flood of experimental gaming studies. By now, several thousands of studies have been published in this area, a number which makes it effectively impossible to consider them all. In a recent review in this area (Liebrand et al. 1992) the following factors were identified as being capable of affecting levels of cooperation in PDG and PDG-like settings:

- monetary payoff structure;
- structure and interpretation of outcome setting;
- availability of communication;
- expectation of others' cooperation;
- others' strategy;
- others' characteristics;
- individual differences;
- group size;
- perceived efficacy;
- perceived risk;
- uncertainty;
- identifiability;
- feelings of responsibility;
- and intergroup structure.

This list of factors surely is not an exhaustive taxonomy. It is more a list of convenient headings for groups of studies known to the authors. Each of the above factors will affect cooperation levels in a complicated way, e.g. obtained levels of cooperation tend not to be a linear function of group size, suggesting that we really need to consider interactions between factors in the above list. As was said before, an integrative framework guiding such research does not exist at the moment and would also be extremely difficult to construct using the traditional approach.

Towards a comprehensive view of research on group decision making

Within science we primarily obtain knowledge of a specific field by reading the published results of theoretical and empirical studies. As long as we can be sure that the body of knowledge for that research area is adequately covered in the literature, there seems to be nothing

wrong with this approach. On second thought, however, this approach may lead to a biased and incomplete perspective.

It might be biased because studies that failed to produce significant results tend not to be published. Speaking from my own experience in the field of social dilemma research, I know of several studies from different laboratories, which failed to find systematic differences for a specific type of framing effect. Almost all of these studies never got published, a few which found such differences were. It would therefore be unlikely that somebody draws a valid conclusion with regard to the effects of this specific manipulation, based on what has been published in this area. The example might seem an exception and insignificant; I am afraid, however, it is not the only one.

The perspective of a research area if based on studying published research findings might be incomplete because, as indicated earlier, it primarily contains the lower levels of knowledge. It concerns mainly, facts, objects, and rules related to knowledge, while heuristic knowledge is underrepresented severely. This kind of high-level knowledge normally is transferred in small workgroups of expert researchers only.

Incomplete or not, the sheer number of published research findings in the area of group decision making renders it almost impossible to identify and integrate the major features of these findings into a theoretical model. What we need here are some inspired guesses or heuristics to search for recognizable patterns. It is therefore desirable to have access to the heuristic knowledge of expert researchers. What we would like to know is how the expert identifies threads in research findings on the basis of which a specific interpretation is chosen out of all possible interpretations. And, in addition, what we also would like to know is how experts go about solving those highly complicated problems with which we confront our undergraduates.

An AI approach

Experts can be used as providers of information, problem-solvers and explainers (Hart 1989). More specific to the area of decision making in situations of outcome interdependency we might use experts to specify recognizable patterns in the voluminous research literature, to identify major issues and problems, and to generate a set of possible solutions to specific problem domains. Such an approach is

a major departure from the way in which we usually develop knowledge of a particular domain.

First, the development will be a task for multiple experts instead of one. From a methodological point of view, the advantage is obvious: it is more reliable to have several estimators instead of one. No matter how good one expert might be, both in terms of declarative and procedural knowledge, the combination of expertise from several individuals is an efficient way to get rid of idiosyncracies.

Eliciting knowledge from experts in the proposed way has the advantage that the resulting knowledge base is highly efficient. The knowledge engineer knows that the acquired knowledge has to be represented in such a way that it can be used in a computer program. As a consequence the knowledge has to be functional with respect to the requirements of the system agreed upon earlier. Irrelevant facts, concepts and rules therefore tend to disappear.

The third advantage lies in the broader domain covered by the experts' knowledge, and more important, the broader domain to which their common heuristic knowledge applies. As a consequence, a group of experts is more flexible in dealing with highly complex situations than a single expert will be. In addition, small group judgment tends to be better than individual judgment and offers the possibility of selecting the optimum solution to a set problem (Greenwell 1988).

On the negative side, it is obvious that the logistics of dealing with several experts are far from being trivial. The use of multiple experts with slightly different fields of expertise requires a study of the overlap of the different domains as well as a thorough study of the ultimate goal of this knowledge elicitation process.

With respect to this goal we have to be modest. One of the prominent criticisms of artificial intelligence is that AI-researchers promised far too much. What we obviously will not obtain in the near future is an expert system for the construction of theoretical frameworks in the area of decision making. However, in my opinion it is realistic to expect significant progress by using knowledge elicitation techniques on a small group of experts in a well-defined domain of decision making.

A concrete example

An example to illustrate the AI approach can be found in the problem of allocating resources. Imagine a supervisor gets a sizeable amount of money as appreciation for the performance of a small group of employees. Our task is to construct a system to advise the supervisor about the impact of possible distributive rules. A computerized system may be cheaper in the long run than hiring experts as consultants.

Given this problem, the next step will be the actual elicitation of knowledge from experts. Broadly speaking this can be done in two ways. The first is 'task oriented' in which the expert is solving the problem and then is observed or interviewed or gets the instruction to think aloud during problem solving. The second is 'solution oriented' in which the expert compares several solutions with the help of structured techniques like Repertory Grids, Multidimensional Scaling and Cluster Analysis. In the present example we follow a method proposed by Butler and Corter (1986) and try to avoid idiosyncrasy in the elicitation method used by combining several of them, i.e. pairwise comparisons, two different scaling techniques and a structured interview (Passchier and Glaudé 1992).

First an American and a Japanese expert in the field of resource allocation were asked to jointly generate several distributive rules (table 1). Next they separately had to compare these rules on a pairwise basis with respect to their similarity. No instruction was given pertaining to the dimension of these similarity ratings. The similarity matrix was then analyzed by a traditional hierarchical cluster analysis

Table 1
Distributive rules generated by the experts.

-
1. *Equal*. Divide equally, each gets the same amount
 2. *Performance*. Divide according to performance. The best performer gets the most, etc.
 3. *Need*. Divide by giving the most needy person the most, etc.
 4. *Future performance*. Divide so as to increase the future performance of the group as much as possible.
 5. *Seniority*. Give the person who has been with the company the longest the most, etc.
 6. *Themselves*. Let the employees themselves decide how to divide the bonus.
 7. *Salary*. Give each employee the same fraction of his or her salary.
 8. *Age*. Give the oldest person the most, etc.
-

and by a Pathfinder network analysis (McDonald and Schwaneveldt 1988). The latter can be considered to be a generalization of a basic tree. Networks, however, are less restrictive than basic trees, because it is possible for a path of node-links to start and end in the same node. It has been shown that this technique yields information which does not show up in the hierarchical cluster analysis (Passchier and Glaudé 1992). Figs. 2a and 2b show the tree and network representation of the eight distributive rules for the American and Japanese expert, respectively. Global inspection shows that both experts differ in their representation. More information about these differences is obtained during the structured interview, the last phase of this approach.

Dendrogram using Complete Linkage

Fig. 2a. Solution for the American expert. Top: Hierarchical clustering solution (SPSS); vertical lines denote joined clusters; the smaller the distance the more similar the rules. Bottom: PATHFINDER network representation; similarity corresponds to lowest number of steps between two rules.

Dendrogram using Complete Linkage

Fig. 2b. Solution for the Japanese expert.

In the last phase, both analyses of the similarity matrix serve as input for a structured interview with the expert in order to elicit the unique features of specific distributive rules or sets of similar distributive rules. The interview itself is structured on the basis of the tree representation. That is, the expert is asked first to indicate the difference between the two most similar distributive rules. Next, the expert needs to specify the difference between, on the one hand, the two most similar rules and, on the other hand, the first distributive rule linking with the first cluster. This process is repeated until the

Table 2

Results of interview with expert A based on hierarchical cluster analysis.

-
1. Seniority vs. Age
Difference relates to length of employment in that organization
 2. Seniority + Age vs. Salary
Salary is external attribute and expresses the summary of the organization's evaluation of employee
 3. Equal vs. Themselves
Distribution by Themselves might create interpersonal conflicts unless they opt for Equal
 4. Equal + Themselves (ET) vs. Seniority + Age + Salary (SAS)
 - Two distinct categories
 - For SAS there is implicit acknowledgement that either of the three is a fair rule, whereas ET does not presume an a priori rule
 5. Need vs. ET + SAS
Need is a third category
 6. Current performance vs. Future performance
Retro- or prospective view
 7. Performance vs. other rules
Performance and Salary should have connection
 - Performance (+ salary) related to motivation and ability
-

whole tree representation has been covered. In this way the linkage pattern generates the stimuli to elicit the unique features of the different distributive rules. Table 2 presents a list of these unique features for the American expert.

For the Pathfinder representation basically the same procedure is used whereby the expert has to indicate differences and resemblances of the various (sets of) rules. Some of this information will be redundant, not all however. First of all, the American expert indicates the existence of the three categories which are nicely clustered in the Pathfinder representation. During the Pathfinder interview the American expert became aware of the very central position of the salary-related rule in his conceptualization. Moreover, the relation between 'performance' and 'salary' is considered to be better represented in the Pathfinder than in the Tree representation.

However, the allocation of resources clearly taps on cross-cultural differences. Especially the Pathfinder interview shows the differences between experts. In terms of the Japanese expert the most easy distributive rule will be to take salary, immediately followed by age and seniority. However, in case all employees are young and the amount of money to be distributed is low, the equality rule might be

the most appropriate. A central point for the Japanese expert is minimalization of conflict. This is illustrated by statements such as, 'the performance rules are bad solutions because workers are a family and you should not differentiate between family members'. To let the employees decide themselves is also a bad solution in Japan because it is 'not done' to talk about bonuses, money or need: 'Samurai do not claim their hunger'.

The type of knowledge system that is represented in figs. 2a and 2b becomes maximally useful when the different rules for allocating resources are linked to different conditions, goals, or constraints. All of the rules that are included in the representation might be appropriate under some set of circumstances, so the final task in the creation of the system is to specify those correspondences. Again, it is the knowledge of the experts that provides the links.

A couple of simple examples can illustrate this component. First, it is obvious that if one were going to use the system to advise a Japanese manager on the allocation of a bonus to his staff, one would rely on the segment of the system that was created by Japanese experts. The rules that are appropriate in the United States may well be inappropriate in Japan and vice versa. Thus, one important type of condition that would have to be specified is the culture context of the application.

The manager's goals are also important. Considerable research has indicated that equal divisions are to be preferred if a major goal of the allocation is to promote cohesion and harmony in the group. However, if the manager wants to use the allocation to maximize the efficiency or the output of the group, than he would be advised to use a performance-based rule. In many real applications the manager might want to accomplish both goals. It would be easy to have the manager specify weights that reflect the relative importance of his objectives, and to use the weights so that an appropriate fraction of the allocation is made by one rule and the remainder by another or others. The final allocation would reflect the mixture of the allocator's goals.

Finally, the system could be designed so that allocations could be recommended, subject to commonly occurring constraints. For instance, one might want to make allocations based on past performance, subject to the constraint that no person gets more than twice as much as another. A different type of constraint might be to set a

maximum and/or minimum allocation, and then to allocate the bonuses in proportion to salary, within these limits.

The rules that are included in the expert's data base can be tested for utility by determining the conditions, goals, or constraints under which a rule would be used. If there are no circumstances under which a rule would be used, then the rule can be safely eliminated from the structure.

In conclusion, it appears that this approach to elicit knowledge from experts is an effective way to make use of the personalized knowledge they have obtained. The overall total amount of time invested in this illustrative example was not more than two days. In addition, the representation of this expert knowledge is sufficiently concrete to use in a knowledge-based computer program. On a more substantial level, even though the example was set up for illustrative purposes, it nicely demonstrates the importance of cultural differences in reward allocation. It is obvious, however, that a more useful implementation also has to specify different conditions, goals and constraints, while it should be based on the knowledge of more than one expert per culture. In my opinion, such an approach would lead to the identification of those facts, concepts and rules that do matter in problem solving. Hence it has considerable value for both theoretical and applied purposes.

Conclusions

It was observed that the area of group decision making currently lacks a comprehensive theoretical framework which integrates the numerous theoretical and empirical research findings. This state of affair probably does not differ from most research areas in the behavioral and social sciences. In addition to the volume of published studies, it might be possible that the published research findings present a distorted view of the existing body of knowledge in a field, and/or represent mainly the lower level types of knowledge.

As a potential solution to this problem we proposed that methods from Artificial Intelligence could be useful to elicit and represent experts' knowledge of facts, concepts, rules and especially their heuristic knowledge in such a way that could be implemented in a computer program. After implementation, the knowledge-based sys-

tem could then assist in decision making in novel situations, based on the synthesis of the experts' knowledge of relevant research findings. The major advantage for the user would be the immediate access to the existing high quality knowledge in that field, without the need to have to study irrelevant material.

More important, however, is the realization that besides the traditional ways to gain knowledge, promising new methods have appeared.

References

- Butler, K.A. and J.E. Corter, 1986. 'Use of psychometric tools for knowledge acquisition: A case study'. In: W.A. Gale (ed.), *Artificial intelligence and statistics*. Reading, MA: Addison Wesley.
- Gammack, J.G., 1987. 'Different techniques and different aspects on declarative knowledge'. In: A.L. Kidd (ed.), *Knowledge acquisition for expert systems*. New York: Plenum Press.
- Greenwell, M., 1988. *Knowledge engineering for expert systems*. Chicester: Ellis Horwood.
- Harmon, P., 1991. The worldwide commercialization of expert systems. Invited address at NAIC/AIT 1991, November 27–38, Free University of Amsterdam.
- Hart, A., 1989. *Knowledge acquisition for expert systems*. London: Kogan Page.
- Liebrand, W.B.G., D.M. Messick and H.A.M. Wilke, 1992. *Social dilemmas; Theoretical issues and research findings*. Oxford: Pergamon Press.
- Luce, R.D. and H. Raiffa, 1957. *Games and decisions: Introduction and critical survey*. London: Wiley.
- McDonald, J.E. and R.W. Schwanefeldt, 1988. 'The application of user knowledge to interface design'. In: R. Guindon (ed.), *Cognitive science and its applications for human-computer interaction*. Hillsdale, NJ: Erlbaum.
- Passchier, P. and J. Glauvé, 1992. *Scaling up knowledge elicitation: Asking the right questions*. Unpublished manuscript, University of Groningen, Social Science Information Technology, Groningen.
- Parsaye, K. and M. Chignell, 1988. *Expert systems for experts*. New York: Wiley.
- Schelling, 1973. Hockey helmets, concealed weapons, and daylight saving; A study of binary choices with externalities. *Journal of Conflict Resolution* 17, 381–428.
- Simon, H.A. and A. Newell, 1958. Heuristic problem solving: The next advance in operations research. *Operations Research*, January–February, 1–10.
- Tuthill, G.S., 1990. *Knowledge engineering: Concepts and practices for knowledge-based systems*. Blue Ridge Summit, PA: TAB Books.