



Book Reviews

Expert Systems: Techniques, Tools, and Applications

Bryan M. Kramer

Expert Systems: Techniques, Tools, and Applications is a collection of research papers covering a range of topics in expert systems and knowledge engineering. The book is edited by Philip Klahr and the late Donald A. Waterman, both of Rand Corporation. The papers are selected from RAND technical reports published from 1977 to 1985. The book is most valuable to people learning knowledge engineering. Four of the papers provide interesting glimpses at the problems involved in transforming knowledge about a domain into computer representations. In addition, the book contains one or two interesting papers for researchers in each of the areas of knowledge acquisition, reasoning with uncertainty, and distributed problem solving.

The book has four sections: introduction, tools, applications, and techniques. Each section begins with an overview that lists the contributions of each of the papers contained therein.

The introductory paper is authored by the editors of the book. It presents a list of the contributions to artificial intelligence of RAND scientists and consultants. For example, Newell, Shaw and Simon were associated with RAND when they did their work on GPS. It is interesting to note how many important researchers have been associated with RAND at one time or another.

The second section contains two papers describing tools for building systems.

The first is a description by Henry A. Sowizral and James R. Kipps of ROSIE, a rule-based expert system shell. The design of ROSIE is based on the assumption that a near-English allows non-programmers to read and understand ROSIE programs. The language is described in great detail, generally using English definitions rather than formal specifications of syntax. Because of the amount of detail, many of the concepts are not very well motivated, and the knowledge representation issues that lead to the design decisions are not discussed.

The second tool is ROSS, an object-

oriented simulation system. The paper is written by David J. McArthur, Philip Klahr and Sanjai Narain. The ROSS system consists of a fairly standard object-oriented language, a predefined object for managing simulation time, some menu-based browsing/command facilities, some EMACS macros for editing ROSS objects, and an interface to a color-graphics system. Programmers are encouraged to make ROSS code English-like by using verbose message templates and by using the abbreviation facility. The paper is valuable as a brief introduction to simulation in an object-oriented programming system.

The third section of the book, applications, contains four papers describing the development of prototype expert systems in four different domains

In the first paper, Donald A. Waterman and Brian M. Jenkins describe an expert system for discovering attributes of terrorist groups given descriptions of their activities. Much of the paper consists of a justification of the expert systems approach and a description of RITA, a rule language developed at RAND in the 1970's. This is followed by a brief description of terrorism, a description of user interaction with the prototype expert system, and examples of the rules and frames used by the expert system.

The second paper in this section is a discussion of the use of an expert system to evaluate claims in product liability disputes. The authors, Donald A. Waterman and Mark A. Peterson, devote much of the paper to a discussion of the settlement process and a list of a large number of unanswered questions. A very small prototype written in ROSIE is then presented. There is then a long discussion of what is yet to be done and of the potential benefits of a completely developed system.

The subject of the third paper is TATR, a prototype expert system for tactical air targeting. The authors are Monti D. Callero, Donald A. Waterman, and James R. Kipps. This is another example of the use of ROSIE. Again, although only a very simple prototype is presented, many of the issues that arise in the knowledge-engineering process are discussed.

The final paper of the applications section is by Philip Klahr, John W. Ellis, William D. Giarla, Sanjai Narain, Edison M. Cesar and Scott R. Turner. The authors present TWIRL, a ROSS-based simulation of tactical warfare. This paper differs from the previous three in that a fairly detailed knowledge base is presented.

The fourth section of the book is entitled "Techniques". It contains four papers describing research in the components of knowledge-based systems. The areas addressed are exemplary programming, knowledge acquisition, uncertain inference, and distributed problem solving.

Exemplary programming is a technique for developing programs from traces of user interactions. The first paper, written by Donald A. Waterman, William S. Faught, Philip Klahr, Stanley J. Rosenchein and Robert B. Wesson, begins with descriptions of two existing exemplary programming systems, EP-1 and EP-2. These are programs that monitor a user performing a series of actions on the computer and create programs for performing similar sequences of actions. In the paper, the shortcomings of EP-1 and EP-2 are described and a detailed proposal for the next generation program is presented. Appendices to the paper contain example sessions with EP-1 and EP-2.

In the second paper of this section, an architecture for knowledge programming is described by Frederick Hayes-Roth, Philip Klahr and David J. Mostow. Knowledge programming is described as a type of knowledge acquisition in which advice expressed in the expert's language is converted into actions in a knowledge-based system. The authors focus on a step they call operationalization; other steps in the process are parsing, interpretation, integration and application. The game of hearts is used as an example; the rules and objectives of the game are transformed into strategies for making moves in the game.

The next paper is written by J. Ross Quinlan. He describes INFERNO, a program embodying a unique scheme for uncertain inference. INFERNO is based on the assumptions that 1) there should be no a priori assumptions of the probabilities of propositions 2) it should be possible to state probabilistic relationships between propositions 3) it should be possible to state information about any subset of the propositions and draw inferences 4) the system should detect inconsistencies. This is a very interesting approach to reasoning with uncertainty.

The last paper describes experiments using different strategies for distributed problem solving. The authors are Randall Steeb, David J. McArthur, Stephanie J. Cammarata, Sanjai Narain, and William D. Giarla. A strategy for distributed problem solving is described as having two parts: a policy for assigning the parts of a task to different agents, and a policy for communication between agents. Two sets of experiments are presented. The first is a simulation of air traffic control. Here the distributed agents are the various planes and the control tower. The second set of experiments is in the domain of remotely piloted fleet control.

In summary, this book is most interesting for the papers in the applications section. Each of these papers describes a domain of expertise and the questions that arise in representing these domains in a knowledge-based system. Each paper gives examples of the final computer representation of the knowledge involved. The

TWIRL paper is particularly interesting in this respect. This section makes the book valuable to anyone learning knowledge engineering, particularly since so few detailed descriptions of complicated systems exist in the literature. In addition, the large number of unanswered questions raised in each paper provide a sobering reminder of how difficult it really is to capture knowledge in a computer system. For builders of expert system tools, the book provides examples of two systems and discussion of some of the questions whose answers will shape future generations of knowledge-based systems.

The Principles and Applications of Decision Analysis and Decision Analysis and Behavioral Research

Jack Breese

Increasingly, artificial intelligence researchers attempting to develop machine encodable theories of knowledge and action are realizing the importance of 1) explicit consideration of the uncertainty (incomplete knowledge) inherent in any real-world domain and 2) the need to handle multiple, possibly conflicting objectives in developing plans of action.

Decision analysis is an engineering discipline which embodies both a theory of decision making under uncertainty and a set of engineering techniques for applying this theory in the real world of multiple objectives. This review examines two important works in decision analysis: *The Principles and Applications of Decision Analysis* edited by Ronald Howard and James Matheson (SDG Publications, Strategic Decisions Group, Menlo Park, California, 1984, \$75.00) and *Decision Analysis and Behavioral Research* by Detlof von Winterfeldt and Ward Edwards (Cambridge University Press, 1986, \$19.95). Though decision analysis typically emphasizes assisting human decision makers in high-stakes complex problems, it has much to offer artificial intelligence both in terms of a theoretical framework for resource allocation in a dynamic and uncertain world, as well

as a set of well-grounded and tested knowledge engineering methods for applying the theory.

At the outset it is important to distinguish between decision theory and decision analysis. Decision theory is a formal system of choice under uncertainty which prescribes that a decision-maker who wishes to abide by a set of axioms that describe a notion of rational and plausibly desirable behavior (von Neumann and Morgenstern, 1947) should choose a course of action that maximizes expected utility. Expected utility is calculated using the standard laws of probability, where the probabilities are subjective, in that they summarize the decision maker's information and beliefs about a proposition. Both of the books reviewed here provide summaries of these basic theoretical underpinnings of decision analysis, as well as some extensions to the theory made by the authors. However, decision theory does not provide any guidance on how to construct a utility function, structure and represent problems, elicit or construct probabilities, and explicitly solve (or search) for recommended courses of action given a particular model or representation. These activities, which are analogous to the tasks of knowledge engineering, representation, and inference in knowledge-based artificial intelligence systems, are the domain of decision analysis. Because the books reviewed here are about decision analysis (as opposed to decision theory), they are devoted primarily to discussion of using rational principles to help people make difficult decisions.

The Principles and Applications of Decision Analysis by Howard and Matheson is a two-volume compilation of tutorials and research papers spanning over 20 years of decision analysis research by Howard and his colleagues at SRI International and Stanford University. Howard and Matheson's primary focus is a view of decision analysis as a procedure for constructing, evaluating, and refining a decision model. This perspective is summarized in the "decision analysis cycle" which is an idealization of a control cycle for decision analysis (Figure 1).

Basis development is analogous to

the knowledge representation task in AI systems. Decision analysis provides a specific taxonomy for this representation: It must be able to capture the decision alternatives, information, and preferences of the decision maker. Deterministic structuring is primarily a means of focussing attention through sensitivity analysis on those factors in the decision situation worthy of probabilistic treatment. Probabilistic evaluation includes development of probabilities, assessment of risk attitude (i.e. the decision maker's preferences for alternatives that have uncertain outcomes), and development of recommended alternatives based on maximization of expected utility. The final phase, basis appraisal, involves assessing the decision model and its results. A critical concept at this stage is the value of information, which provides a measure of whether it is worthwhile to gather additional information. In particular, information is valuable only to the extent that what one learns might change a decision. The cycle iterates as long as it is necessary to revise the model or there are opportunities to profitably collect information (i.e. reduce uncertainty). The exposition of the view of decision analysis illustrated in Figure 1. is enhanced by many examples in the introductory articles of the two-volume set, as well as published works reporting the results and findings of real applications of decision analysis in industry and government.

The volumes are organized into six major areas- Introduction and Overview, Applications, Health and Safety, Professional Practice, Technical Contributions, and Contributions from Psychology. Artificial intelligence researchers perhaps can benefit most from the sections on professional practice which contain papers on probability assessment, risk preference, and influence diagrams. Probability assessment is the knowledge engineering task of helping an expert or decision maker express his information about an event as a probability. Risk preference refers to a theory describing an individual's propensity for engaging in risky actions. Influence diagrams are a representation of decision problems under uncertainty

which allows one to represent decision and state variables, information flows, and value and utility relations in a directed graph. This representation has been adapted by artificial intelligence researchers for doing strictly probabilistic (no decisions or values) reasoning (Pearl, 1986).

The section on Technical Contributions also contains several articles likely to be of interest to the AI community—for example "Information Value Theory", "Value of Information Lotteries", and "Proximal Decision Analysis" by Howard and "The Economic Value of Analysis and Computation" by Matheson. These articles, most of which appeared originally in systems and management science journals, are technical works which very clearly and concisely develop and demonstrate key ideas in decision analysis. "Proximal Decision Analysis," for example, shows how large decision problems can be analyzed with approximate methods by using derivative information and sensitivity analysis.

Since it is a compilation of papers, the Howard and Matheson volumes contain very concise, insightful treatments of the basic theoretical underpinnings of decision analysis such as basic probability, utility theory, and the decision analysis cycle, as well as focused discussion of several more advanced topics including value of information, probability assessment, sensitivity analysis, and others. However, the volumes can be somewhat frustrating to use because much of the basic material on decision analysis is repeated in several articles; while some of the important technical material is difficult to find. A major deficiency for an uninformed reader is the lack of an index to the topics in the various papers.

In addition, there are some gaps in the exposition. For example, there is no complete discussion of Howard's distinctions between value (risk-free scoring of alternative states) and utility (scoring of alternative states under uncertainty) distinctions and associated theoretical treatments. Except in extreme life and death situations, Howard generally reduces each attribute of a multiattribute decision problem to a single numeraire —

money—in order to facilitate value function assessment. Utility is assessed by applying a transformation encoding risk preference to the value function. This an important heuristic that Howard has found useful for the business, strategic, and health and safety problems which have been the subject of much of his professional practice and research, however the formal basis and practical efficacy of this perspective is not discussed in these volumes except by example. This method for encoding values and preferences is an additional facet which distinguishes the Howard-Matheson school of decision analysis from other perspectives. *Decision Analysis and Behavioral Research* represents one of these other perspectives. Howard and Matheson are engineers in that they view decision analysis fundamentally as a technology for assisting people making difficult decisions. They regard contributions from psychology peripherally as corroborating that people make errors in reasoning and decision making in uncertain situations and secondarily as information for improving assessment procedures. *Decision Analysis and Behavioral Research* is a collaborative effort of von Winterfeldt, a psychologist, and Edwards, a decision analyst, originally trained as a psychologist. This collaboration results in a book which is grounded in its discussion of the prescriptive facets of decision analysis with respect to its behavioral psychology precursors and implications. The relevance of psychological research to both the theory and practice of decision analysis is integral to this volume.

This book contrasts with the Howard and Matheson work in that it is a unified and complete treatment of a wide range of topics in decision analysis. The technical material is in most cases first presented informally by way of example, and then followed by a more rigorous, mathematical treatment. Overall, the emphasis is on verbal description and motivation wherever possible. In addition to several chapters on basic decision analysis and probability, there are several sections on more advanced topics, including a fairly complete discussion of probability measurement from a

personal or subjective point of view. This is coupled with a discussion of the insights that psychological literature provides for careful and accurate assessment of probabilities. Chapters on inference include a discussion of Bayesian statistics (use of Baye's rule for subjective probability in statistical inference). The more general inference problem is also discussed, including an extended example of legal decision making and reasoning under uncertainty. This discussion should be of particular interest to AI researchers, since it is an attempt by Bayesian analysts to tackle complex real-world inference problems with probability machinery.

Von Winterfeldt and Edwards contend that decisions are often difficult due to lack of clarity about the decision maker's values and preferences about the outcomes resulting from alternative courses of action. Therefore, they emphasize careful structuring and assessment of preferences as being central to the process of doing decision analysis. The book encapsulates and extends previous work on value and utility measurement (Keeney and Raiffa, 1976) for both single attribute and multiattribute cases. One of the major research results of this volume is a rejection of the distinction between notions of value and utility. This distinction has played a major role in most other work in decision theory and decision analysis, including that of Howard and Matheson. Von Winterfeldt and Edwards develop a set of arguments, on both theoretical and practical grounds, for rejection of the distinction, while at the same time recapitulating previous work in multiattribute utility theory. In these chapters, as well as those on probability, the authors emphasize where 1) the literature indicates that the decision theoretic constructs are inconsistent with actual behavior in experimental situations and 2) how knowledge of the nature of these departures can be accounted for in eliciting and constructing utility scales. The book contains some additional interesting chapters and sections. The chapter on sensitivity analysis addresses the sensitivity of decision-analytic models to changes in input assumptions and errors of mea-




Figure 1. The Decision Analysis Cycle

surement. A discussion of the pitfalls of decision analysis is primarily a recognition of how a decision analysis can be subverted in the real world due to hidden agendas, recalcitrance of participants, and multiple organizational objectives. Cognitive illusions, i.e. the errors of judgement under uncertainty that have been reported in the psychological literature, are reinterpreted. These topics, along with a summary of the history of decision analysis, results in a book which covers the field of decision analysis in both breadth and depth from a psychological perspective.

The Howard and Matheson compilation and the von Winterfeldt and Edwards book both cover the same basic topics though from somewhat different perspectives. Howard and Matheson emphasize modelling, sensitivity analysis, and value of information while von Winterfeldt and Edwards stress the structure of preferences, measurements, and relationships to behavioral research. Both books provide a comprehensive view of decision analysis and its use in helping individuals and organizations in the process of gaining insight by clarifying objectives, forming models, and generating recommendations.

What is the significance of decision analysis for artificial intelligence? The field of decision analysis has grown as a consulting and engineering activity in the analysis of high stakes, complex, policy or strategic decisions with considerable uncertainty. These problems are typically one-shot or occasional decisions where the expense and effort of a decision analysis is justified by the importance and complexity of the situation. Over approximately the same period, artificial intelligence researchers have been attempt-

ing to imbue computer systems with abilities to form goals, manage beliefs, form and execute plans, and allocate resources in an uncertain and dynamic world. The nature of the problems addressed by AI are typically of lower stakes (common-sense reasoning, planning, diagnosis and trouble-shooting, etc.) and occur with greater frequency than those to which the decision analytic machinery has been applied. However regardless of the stakes or frequency, the need to couple of decision making (i.e. resource allocation) and uncertainty (lack of information) is important across this entire spectrum.

The question of whether some of the techniques of decision analysis for humans can be fruitfully applied to AI resource allocation problems is just beginning to be explored. In this light, decision analysis is a potentially rich field for AI researchers for theoretical concepts, knowledge representations, inference techniques, and knowledge engineering methodologies. Both the works described here are admirable compilations of the current state of the art and can provide both AI researchers and practitioners with a new perspective and set of methods for problem solving.

References

- Keeney, R. L. and H. Raiffa, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, New York: John Wiley and Sons, 1976
- Pearl, J. "Fusion, Propagation, and Structuring in Belief Networks," *Artificial Intelligence*, (29:3), pp. 241-288, September 1986.
- Von Neumann, J. and O. Morgenstern, *Theory of Games and Economic Behavior*, Princeton, New Jersey: Princeton University Press, 1947