AI Magazine Volume 19 Number 3 (1998) (© AAAI)

Interactive and Mixed-Initiative Decision-Theoretic Systems

Peter Haddawy and Steve Hanks

■ The American Association for Artificial Intelligence Spring Symposium on Interactive and Mixed-Initiative Decision-Theoretic Systems was held at Stanford University from 23–25 March 1998. The symposium attracted approximately 30 researchers from around the world. Topics discussed included incremental model construction, user interaction, explanation generation, and applications.

Decision-theoretic techniques were originally developed to help make careful choices in high-stakes situations such as making large business investments, planning military strategy, and choosing among medical treatment alternatives. Researchers in the uncertainty in AI community are asking whether these techniques can also be applied to help make high-quality decisions in more commonplace situations where the stakes are lower.

Decision theory is an attractive framework for building interactive systems for decision making or decision support, but a traditional decision-theoretic analysis requires both a probability model and a utility model, and it is typically time consuming and tedious to elicit either one. This overhead might not be justified by the importance of the problem being solved, especially if the elicitation cost cannot be amortized over many problem-solving episodes.

Often a single problem can be solved without eliciting a complete model in advance. For example, an automated travel agent would not need information about all of a user's travel preferences to build a single itinerary. Because it is usually impossible to ascertain ahead of time exactly what preference information will and will not be relevant to solving a particular problem, there is a need to interleave the elicitation of preference information with the problem-solving process itself.

Fortunately, the richness of the decision-theoretic framework provides valuable flexibility in problem representation. Unimportant portions of a problem space can be represented using coarse preference information or omitted altogether, but more important parts of the decision space can be represented more precisely. Symposium participants presented several techniques for representing and computing with partial or abstract models. Various techniques were presented for eliciting the decision model incrementally in conjunction with the problem-solving process. Well-established techniques from decision analysis, including sensitivity analysis and value of information calculation, were also discussed in the context of incremental model elicitation. Finally, the importance of self-explanatory systems was emphasized because the user needs to understand the impact of his/her communicated preferences and their role in the problem-solving process.

Decision theory is an attractive framework for building interactive systems for decision making or decision support, but a traditional decision-theoretic analysis requires both a probability model and a utility model

Peter Haddawy received his Ph.D. in computer science in 1991 from the University of Illinois at Urbana. Since then, he has been teaching at the University of Wisconsin at Milwaukee, where he is associate professor and director of the Decision Systems and Artificial Intelligence Lab. He is currently on leave at Assumption University in Bangkok, Thailand, where he is acting director of the Intelligence Systems Lab. His research interests include representations and algorithms for planning and decision making under uncertainty, use of abstraction in decision-theoretic problem solving, incremental problem-solving techniques, and medical applications. His e-mail address is haddawy@cs.uwm.edu.

Steve Hanks is manager of the Adaptive Systems Group at Harlequin Inc. The group applies AI-inspired techniques to a variety of products in areas such as digital publishing and printing, information retrieval, and memory management. Prior to joining Harlequin, Hanks was associate professor of computer science at the University of Washington. He completed his Ph.D. at Yale University under the supervision of Drew McDermott. His research has touched on a variety of topics related to decision-theoretic reasoning, including probabilistic planning, probabilistic temporal reasoning, and the incremental assessment of utility models. His e-mail address is hanks@harlequin.com.