DECISION ANALYSIS
Vol. 2, No. 4, December 2005, pp. 183–184
ISSN 1545-8490 | EISSN 1545-8504 | 05 | 0204 | 0183

informs

DOI 10.1287/deca.1050.0061

From the Guest Editor . . . New Contributions and Reflections on Graph-Based Representations for Decision Analysis

Eric Horvitz, Guest Editor
Microsoft Research, One Microsoft Way, Redmond, Washington 98052, horvitz@microsoft.com

In this second volume of the special issue of *Decision Analysis* on graph-based representations for decision analysis, we present two articles and four brief invited *perspective* pieces, capturing personal reflections on the theoretical and practical influences of probabilistic graphical models in several realms.

Key words: decision analysis; influence diagrams; Bayesian networks; graphical models; editorial

With this second volume, we complete the special issue of *Decision Analysis* on graph-based representations for decision analysis. In the first volume (September 2005), we published Howard and Matheson's original article on influence diagrams, along with a retrospective by the authors, and two new manuscripts. The current volume includes two additional articles followed by several shorter invited perspectives on the influence of graphical models.

In the first article, Ali Abbas and Ronald Howard present a class of utility models that they refer to as *attribute dominance* utility functions. They describe how methods for assessing joint probability distributions can be used to assess these utility functions and demonstrate the generality of their approach by showing how any utility function can be reformulated as an attribute dominance utility function. The authors introduce an intriguing graphical representation of attribute dominance utility functions, named *utility diagrams,* and present a "utility inference" procedure that is analogous to probabilistic inference.

The second paper, by Apiruk Detwarasiti and Ross Shachter, moves beyond traditional decision analyses, which consider the beliefs and preferences of a single principle agent, to tackle *team decision making* challenges. The authors use graphical models to represent decisions that involve groups of people, such as actions that are taken by decision makers who are not collocated. They examine the class of team decision problems where team members are assumed to have common beliefs and preferences but cannot explicitly share information, and map such team decision-making problems to actions by a single person with imperfect recall. They introduce influence diagram representations of the team decision problems and show how the graphical models can elucidate opportunities for simplification and optimization.

The constellation of short perspectives following the longer contributions serves as an interesting ensemble of reflections on the history and influences of probabilistic and decision-theoretic graphical models. The authors of the perspectives share their thoughts on how graphical representations have shaped scholarship and practice in several realms, including insights about accomplishments to date and challenges ahead.

Graphical probabilistic models, including influence diagrams and Bayesian networks, have played a central role in the evolution of research in artificial intelligence (AI), where investigators have explored theoretical and practical problems with the automation of learning and reasoning. Historical analyses and reviews of the current state of research on graphical models in AI, statistics, and decision analysis reveal complementary cross-discipline efforts and interactions, with both shared and distinct goals and motivations.

Craig Boutilier reflects about the impact of influence diagrams in AI. He reviews how graphical models, and decision-theoretic concepts more broadly, have sculpted in a very significant manner research directions and approaches in the AI community. Judea Pearl shares personal reflections about the relationships between his work on Bayesian networks and the preceding work on graphical models in the decision analysis community. Readers will likely find interesting Pearl's recollections of his early meeting with Howard, Matheson, and colleagues. His comments underscore how differences in fundamental passions, goals, and curiosity led to different stresses and priorities in work on graphical models. Pearl, and others in the AI community who have pursued the dream of understanding intelligence and of building automated reasoning systems, investigated graphical probabilistic models as a promising representational fabric for automated learning and inference. These goals motivated such efforts as the derivation of graph-theoretical results about the specification of independence and the formulation and testing of several exact and approximate algorithms for propagating beliefs in graphical models containing large numbers of variables. In earlier work, Howard, Matheson, and colleagues in the decision analysis community pursued influence diagrams as a means for enhancing the practice of decision analysis, with a focus on easing assessment and entailment. Pearl noticed with interest, and some surprise, the differences in thrusts and goals in the initial meeting he had with Howard and Matheson on graphical models.

Over the last two decades, influence diagrams have been used by practitioners of decision analysis for tackling hard problems in numerous domains. Dennis Buede shares his personal experiences with using influence diagrams for tackling challenging decision problems that he has encountered in the course of his work as a professional decision analyst. He reports on the value of using influence diagrams in consultations he has performed in working with government agencies, including challenges defined within the intelligence community.

Stephen Pauker and John Wong share their thoughts on the role of influence diagrams in medical decision making. Decision analysis has played a starring role in evidence-based medicine, providing insights to clinicians, patients, and policymakers about challenging, real-world health-care decision problems. A thriving community of medical decision making (MDM) researchers came into existence over 25 years ago and assembles each year for the annual MDM meeting. Pauker and Wong describe the continuing dominance of decision trees in medical decision making and reflect about the relatively poor penetration of influence diagrams in their community. The authors' reflections are relevant for professional decision analysis in domains beyond health care, and their insights highlight challenges and frame opportunities for future innovation.

The special volumes of *Decision Analysis* have provided a venue for recognizing formative work on influence diagrams and communicating new results at the frontier of research. They have also provided a place for sharing reflections from several communities about progress and directions on graphical representations in decision making. Given the growing density of research on graphical probabilistic and decision-theoretic representations—and the expected stream of future publications on graph-based methods—I suspect that people will one day look back upon our decision to prepare "special" volumes on graph-based methods with interest and curiosity.

It has been an honor to serve as guest editor of the two volumes.