DECISION ANALYSIS
Vol. 2, No. 3, September 2005, pp. 125–126
ISSN 1545-8490 | EISSN 1545-8504 | 05 | 0203 | 0125

informs

DOI 10.1287/deca.1050.0051


# From the Guest Editor . . . Special Issue on Graph-Based Representations

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

Graphical representations have played a central role in decision analysis. Although decision trees remain popular, more general graphical languages can be used to encode relationships among variables of a decision basis. Influence diagrams, introduced in the late 1970s, provide a general graphical representation for decision analysis. We publish, in this first of two volumes of a special issue of *Decision Analysis*, the original manuscript on influence diagrams, along with a retrospective by the authors. We also include two new articles focusing respectively on the use of graphical models for group coordination challenges and on handling the unintended side effects of interventions.

*Key words*: decision analysis; influence diagrams; graphical models; editorial

Graphical representations of beliefs and preferences have played a central role in decision analysis. For at least as far back as the time of Blaise Pascal, decision analysts have represented actions, uncertainties, and preferences with branching trees capturing paths through a conceived space of outcomes. Although decision trees remain popular, more general graphical representations have come to the fore. Today, decision analysts and computer scientists have access to a variety of directed and undirected graphical representations for encoding relevance, independence, and sequence, including influence diagrams, Bayesian networks, dependency networks, Markov networks, and CP-nets. Graphical models have proven to be valuable for eliciting influences, beliefs, and preferences in a natural and comfortable manner, and for providing a computational substrate for efficient inference.

The introduction of influence diagrams by Ronald Howard and James Matheson in the late 1970s was a landmark contribution, providing a general graphical representation for decision analysis. Influence diagrams were forged amidst pressures to solve increasingly complex decision problems and the growing availability of computers that could represent and propagate assertions about relationships among variables of a decision basis. Howard, Matheson, and colleagues on their team, recognized that computers could be harnessed to free decision analysts from tedious bookkeeping, while providing both professionals and lay people with a more human-oriented graphical language for specifying, refining, and interacting with representations of decision problems.

It is said often that decision analysis is more about generating insights than about simply identifying actions with maximum expected utility. Influence diagrams shine most brightly in the communication channel they open for nonexpert principal agents, engaging them, reducing the burden of assessment, and, ultimately, in providing insights via an intuitive, yet rigorous method for encoding and reflecting about decision challenges.

Beyond their use in decision analysis, influence diagrams have proven valuable in other ways, including their use as a tool for crisp communication among collaborating experts, for expressing and solving technical decision-analytic challenges in a graph-theoretic manner, and for serving both as a clarifying lens and as an explicit computational representation in artificial intelligence research. On the latter, influence diagrams are well-known within the artificial intelligence community and have been used in theoretical studies, such as efforts pursuing an understanding of causation and causal reasoning, and in engineering efforts

Decision Analysis 2(3), pp. 125–126, © 2005 INFORMS

focused on building and fielding automated decision-making systems.

Howard and Matheson's initial article on influence diagrams expressed with atypical clarity the utility and beauty of the representation. Unfortunately, the manuscript was never published in the archival literature. Those interested in seeing the original work have had to expend varying amounts of effort to track down the 1981 technical report.

We decided to publish Howard and Matheson's original article in archival form, along with a retrospective by the authors, to anchor a two-volume special issue of the journal on graph-based representations for decision analysis. Beyond the article and retrospective, this first volume includes two new articles. In the next volume of Decision Analysis, we will publish two additional articles and several shorter perspective pieces on the influences of influence diagrams, composed by leaders from several disciplines.

On the new articles in this volume, David Pennock and Michael Wellman explore the value of using graphical models to probe belief aggregation and risk sharing—two fundamental problems that come under the herald of *group coordination*. Research on forming consensus beliefs and preferences by combining the beliefs and preferences of groups of people has been fraught with negative results that come in the form of paradoxes and impossibility theorems. Pennock and Wellman make progress in this difficult space, harnessing graphical models to represent the aggregate probability distribution for a group of individuals. Although their analysis highlights new limitations and impossibility results, they also show how graphical models can be used to identify restrictive conditions that allow for consistent aggregations. The paper provides interesting background on group coordination problems and shows how graphical models can be used to build insights and generate results.

The article by David Matheson and James Matheson might be characterized as "the Heisenberg uncertainty principle meets decision analysis." Motivated by their experiences with decision-analysis consultations, the authors investigate challenges with modeling situations where observational or control actions may have unintended, and typically unmodeled, influences on multiple aspects of the decision basis. They introduce the concept of *purity* of interventions and describe extensions to basic influence diagrams that account for side effects of interventions. Then, they provide an approach to estimating the value of interventions in advance of expending the effort to extend decision models to incorporate the side effects.

It has been an honor to serve as the Guest Editor of these volumes. I hope that the collection of articles and commentary in these volumes will serve to appropriately honor the original work on influence diagrams, highlight current research, and ultimately stimulate new research efforts. I thank the many insightful reviewers who diligently read submissions and revisions and provided valuable comments to the authors and editors. I would like to also thank the Editors in Chief of *Decision Analysis*, Robert Clemen and Don Kleinmuntz, for their enthusiastic support of the special issues. I want to especially recognize Bob, who put a great deal of effort into the special issue as a collaborator on the project.