DECISION ANALYSIS
Vol. 2, No. 4, December 2005, pp. 229–231
ISSN 1545-8490 | EISSN 1545-8504 | 05 | 0204 | 0229

DOI 10.1287/deca.1050.0054

The Influence of Influence Diagrams on Artificial Intelligence

Craig Boutilier
Department of Computer Science, University of Toronto, Toronto, Ontario, M5S 3G5 Canada, cebly@cs.toronto.edu

Howard and Matheson's article "Influence Diagrams" has had a substantial impact on research in artificial intelligence (AI). In this perspective, I briefly discuss the importance of influence diagrams as a model for decision making under uncertainty in the AI research community; but I also identify some of the less direct, but no less important, influences this work has had on the field.

Key words: influence diagrams; decision theory; artificial intelligence; value of information; graphical models; perspective, the focus on graphical modeling research
History: Received on November 14, 2005. Accepted by Eric Horvitz on November 23, 2005, without revision.

Howard and Matheson's (1984/2005) "Influence Diagrams" has had a profound impact on developments in artificial intelligence. Some of these influences have been quite direct; others are more indirect, but in many ways, more substantial. The paper itself is representative of developments that had been underway throughout the 1970s, and crystallized the state of the art in graphical models for decision making at that time. "Influence Diagrams" is remarkable for its clarity of exposition and makes a compelling case for the use of graphical models to capture the inherent logic of a decision problem. The influence of this paper can then be seen as emblematic of the impact of influence diagrams as a specific formalism, and more broadly of graphical models for decision making, on artificial intelligence (AI).

The most direct impact on AI is, of course, rather obvious. Influence diagrams (in various guises and under various names) have become a standard modeling tool for decision making under uncertainty within AI. Medical diagnosis offered one of the earliest application areas of influence diagrams in AI, lying at the intersection of AI and and medical informatics/decision making—see Horvitz et al. (1988) for an early survey. This remains an area where influence diagrams find widespread use. However, as part of the standard AI toolkit, they have been used in applications and research domains as varied as diagnosis (medical and otherwise), control systems, computer vision (Binford and Levitt 2003), dialog management, user interface design, multiagent systems, and game theory (Koller and Milch 2003), to name but a few.

Another reasonably direct impact of "Influence Diagrams" derives from its role in the development of graphical models for probabilistic modeling and inference. Probabilistic graphical models, as a general concept, have had a significant impact on almost all aspects of AI. Such models sit at the core of much modern work in inference and reasoning, planning and decision making, machine learning, pattern recognition and computer vision, computational linguistics and information retrieval, and many other research areas. Along with the work of Pearl, Lauritzen, Spiegelhalter, and others, Howard and Matheson's (1984/2005) paper shaped an entire discipline through its compelling vision: Intuitive qualitative models capturing the logic of probabilistic and informational dependence can aid in the construction, understanding and computational solution of complex quantitative inference, learning and decision problems.

Somewhat less direct, but at least as profound as the influences described above, is the gradual change in culture within AI: Decision-theoretic models have not only come to play a central role in AI, but the *language of decision theory and decision-theoretic concepts* now pervade most areas of AI. This shift was directly precipitated by "Influence Diagrams." In the

1970s and early 1980s, probabilistic and especially decision-theoretic models were largely out of favor within AI (certainly within the dominant expert-system paradigms of the day), due in part to "a perception that decision-theoretic approaches were hopelessly intractable and were inadequate for expressing the rich structure of human knowledge" Horvitz et al. (1988, p. 257). Even as (rule-based) expert systems fell out of favor, the logicist paradigm continued to dominate AI through the early 1990s. However, the introduction of graphical models to AI in the early 1980s, especially Bayesian networks and influence diagrams, precipitated a sea change.

"Influence Diagrams" demonstrated that the language of decision theory was not only rich enough to capture the intricacies of complex decision problems faced throughout AI; but with a suitable decomposition of a problem into the representational levels of *relation*, *function*, and *number* (to use Howard and Matheson's 1984/2005 terminology), decision theory provided *the most natural means* of specifying such decision problems. Developments in sophisticated inference methods for Bayesian networks and solution techniques for influence diagrams also demonstrated the computational power of exploiting the structure of independence at both the relational and functional levels. It is the naturalness of the graphical formalism (capturing both probabilistic and informational dependencies and independencies) and subsequent computational developments that ultimately led to the prevalence of decision-theoretic models.

How do decision-theoretic models manifest themselves in AI? In some cases, even when the influence diagram model is not used directly, its characteristics are evident. For example, since the mid-1990s, fully and partially observable Markov decision processes (MDPs) have become a central model for planning and decision making within AI. MDPs had been studied extensively in decision analysis, economics and operations research since the 1950s (Aström 1965, Bellman 1957, Howard 1960, Shapley 1953). Tatman and Shachter (1990) demonstrated how influence diagrams could be used to model finite-horizon MDPs directly; subsequent work on exploiting structured graphical models computationally has led to the widespread acceptance of MDPs as a tool for AI planning and related problems (see Boutilier et al. 1999 for a survey of this area, and Guestrin et al. 2004 for a recent example). The considerable body of work on using graphical models to specify and solve MDPs traces its origins directly to Howard and Matheson's (1984/2005) article, which in turn has led to the widespread application of MDP models in a wide variety of AI applications and research areas.

The indirect impact is evident in every area of AI, where decision-theoretic modeling is now ingrained in the research enterprise. In areas ranging from computer vision and computational linguistics, to machine learning and planning, probabilistic models and cost or utility functions (often using graphical models) are ubiquitous. To take one example within machine learning, both active learning and reinforcement learning are areas in which not just decision theory, but the concept of *value of information* plays a vital role.[1] Yet another impact of "Influence Diagrams" is the fact that it is one of the earliest papers to expose the value of information concept to the AI community.

Like the three levels of specification articulated by Howard and Matheson (1984/2005), the article "Influence Diagrams" has had three "levels of influence" on work in AI: direct utilization of influence diagrams and their variants, the broader adoption of graphical models within AI, and the ubiquity of decision-theoretic modeling and conceptualization throughout AI. All of these have shaped, and continue to shape, the discipline in dramatic ways.

References

Aström, K. J. 1965. Optimal control of Markov decision processes with incomplete state estimation. *J. Math. Anal. Appl.* **10** 174–205.

Bellman, Richard E. 1957. *Dynamic Programming*. Princeton University Press, Princeton, NJ.

Binford, Thomas O., Tod S. Levitt. 2003. Evidential reasoning for object recognition. *IEEE Trans. Pattern Anal. Machine Intelligence* **25**(7) 837–851.

Boutilier, Craig, Thomas Dean, Steve Hanks. 1999. Decision theoretic planning: Structural assumptions and computational leverage. *J. Artificial Intelligence Res.* **11** 1–94.

[1] The exploration-exploitation notion in reinforcement learning, for instance, while typically approached heuristically within AI, has been treated in a more principled, Bayesian fashion by explicit treatment in terms of value of information. While this principled approach is found in the early literature on the topic (e.g., Satia and Lave 1973), recent work in AI has exploited graphical structure in the approximation of exploration policies (e.g., Dearden 2000).

Dearden, Richard. 2000. Learning and planning in structured worlds. Ph.D. thesis, University of British Columbia, Vancouver, Canada.

Guestrin, Carlos, Milos Hauskrecht, Branislav Kveton. 2004. Solving factored mdps with continuous and discrete variables. *Proc. Twentieth Conf. Uncertainty in Artificial Intelligence*, Banff, Alberta, Canada, 235–242.

Horvitz, Eric, John S. Breese, Max Henrion. 1988. Decision theory in expert systems and artificial intelligence. *Internat. J. Approximate Reasoning* **2** 247–302.

Howard, Ronald A. 1960. *Dynamic Programming and Markov Processes*. MIT Press, Cambridge, MA.

Howard, Ronald A., James E. Matheson. 1984/2005. Influence diagrams. R. A. Howard, J. E. Matheson, eds. *Readings on the Principles and Applications of Decision Analysis II.* Strategic Decision Group, Menlo Park, CA, 719–762. Reprinted, *Decision Anal.* **2**(3) 127–143.

Koller, Daphne, Brian Milch. 2003. Multi-agent influence diagrams for representing and solving games. *Games Econom. Behavior* **45**(1) 181–221.

Satia, J. K., R. E Lave. 1973. Markov decision processes with uncertain transition probabilities. *Oper. Res.* **21** 728–740.

Shapley, Lloyd S. 1953. Stochastic games. *Proc. National Academy Sci.* **39** 327–332.

Tatman, Joseph A., Ross D. Shachter. 1990. Dynamic programming and influence diagrams. *IEEE Trans. Systems, Man Cybernetics* **20**(2) 365–379.