

Influence Diagrams—Historical and Personal Perspectives

Judea Pearl

Cognitive Systems Laboratory, Department of Computer Science, University of California,
Los Angeles, California 90095, judea@cs.ucla.edu

The usefulness of graphical models in reasoning and decision making stems from facilitating four main computational features: (1) modular representation of probabilities, (2) systematic construction methods, (3) explicit encoding of independencies, and (4) efficient inference procedures. This note explains why the original introduction of influence diagrams, lacking formal underpinning of these features, has had only mild influence on automated reasoning research, and how Bayesian belief networks, which were formulated and defined directly by these features, became the focus of graphical modeling research.

Key words: Bayesian networks; probabilistic influence; causal modeling

History: Received on November 23, 2005. Accepted by Eric Horvitz on December 9, 2005, without revision.

Influence Diagrams (IDs) command a unique position in the history of graphical models. On the one hand, they can be seen as an extension of path diagrams (Wright 1921) a tradition in which qualitative domain knowledge and assumptions are expressed (often unwittingly) in graphical form, while quantitative statistical information is obtained from empirical data. On the other hand, IDs can also be viewed as informal precursors to belief networks (later called Bayesian networks), which currently serve as the main computational tool for automated reasoning.

Oddly, however, the introduction of influence diagrams have had only mild influence on both the path diagrams and the automated reasoning communities; this note attempts to explain why.

From path-diagrammatic perspectives, IDs permit researchers to break away from the confines of functional linear models and express any relationships whatsoever between variables of interest. However, econometricians and social scientists, the main users of path diagrams, were not ready to put this extension into use, since IDs, as presented in (Howard and Matheson 1984/2005) required subjective assessment of conditional probabilities, and statistically trained researchers were conditioned to mistrust and avoid such assessments at all cost. In contrast, these researchers were accustomed to infer the

parameters of path diagrams from the data itself, sometimes even in the presence of unmeasured confounders. The respective learning techniques for nonlinear IDs, which correspond to identification techniques in causal Bayesian networks (Pearl 2000), were tackled only in the mid 1990s, more than a decade after the inception of IDs.

From an automated reasoning perspective, IDs came into being when the field was preoccupied with a fierce debate between rule-based and probabilistic inference systems—the former offering computational efficiency and the latter providing coherence and theoretical underpinning (Pearl 1993). While the representational economy offered by IDs could well have assisted the probabilistic side in that debate—it did not, because IDs were not accompanied with the computational tools that would render them competitive to rule-based systems. Such tools were uncovered later on, once the conditional independence semantics of Bayesian networks was established (i.e., through the d -separation criterion)

There was little interaction in the early 1980s between AI and decision analysis researchers. When I presented belief propagation on trees (Pearl 1982) before Ron Howard's group at Stanford, the audience could not understand why I emphasized computational features such as autonomous, asynchronous

propagation when the most difficult task in decision problems was that of knowledge elicitation. On my part, I could not appreciate why the Stanford group was excited about IDs when they did not compute anything interesting with those diagrams, and did not even use the diagram to infer independencies that were not already assumed in its construction.

Motivated by problems of parallel processing in visual perception, I was influenced primarily by works on hierarchical inference (Kelly and Barklay 1973) and reading comprehension (Rumelhart and McClelland 1982), and came to view graphical representations from a different perspective. I understood a graph to be an approximate representation of the conditional independence relations that are embedded in a given distribution (Pearl 1986). A DAG that captures a maximal number of such independencies was defined as a *Bayesian network* of P (Pearl 1988, p. 119). Note that *causal* Bayesian networks represent a marked departure from this view; they should be understood as collections of *stable* mechanisms in the domain, bombarded by independent perturbations of those mechanisms (Pearl 2000, p. 21).

In retrospect, I believe there were four main factors that prevented the formulation of Howard and Matheson (1984/2005) from having a more significant impact on automated reasoning research in the 1980s.

First, the formulation was laden with informal, nontechnical jargon that left ample room for ambiguities. For example, “[An arrow pointing from A to B] means that the outcome of A can influence the probabilities associated with B ”¹ (p. 103), or “the states of information upon which independence assertions are made” (p. 131), or “An arrow . . . may be reversed provided that all probability assignments are based on the same set of information” (p. 132), or “the probability assignment of variable g is in principle conditioned on all variables . . .” (p. 135). Today, armed with the notions of conditional independence, I -maps, minimal I -maps and colliders, one can perhaps find consistent interpretations of these informal phrases. Yet, in 1981, they deterred researchers from even trying and drove them, this writer included, to first seek

a basic, formal understanding of the relationships between graphs and probabilities.

Second, the presentation of IDs was not constructional. In other words, although the defining equations

$$\{x \mid N_x, E\} = \{x \mid D_x, E\} \text{ for all } x$$

turn out to be adequate for testing whether a given diagram is an I -map of a given distribution P , they do not permit a step by step construction of a graph that is an I -map of P . The reason is that when the set of parents D_x is selected (using the equations above), we do not know which variables are in the set N_x and which are descendants of x . A constructive definition should allow us to select x 's parents by considering only x 's predecessors in a given construction order. The validity and equivalence of ordered constructions follow from the graphoid axioms (Pearl 1988, p. 120).

Third, the initial presentation of IDs was not interpretational, namely, it did not instruct us how to read from the diagram conditional independencies that are implied by the defining equations above, although not by any one such equation in isolation. For example, all variables that are d -separated from x by a subset of x 's parents represent a valid independency that is not explicitly recognized by the defining equation for x . The d -separation criterion (Pearl 1986) provides the needed interpretation of a diagram in terms of the conditional independencies it imposes.

Finally, the introduction of IDs was not accompanied with computational procedures that utilize the structure of the diagram to facilitate probabilistic inferences. Such procedures were developed later, mostly in the framework of Bayes networks,² for which the semantics of conditional independencies was made explicit. The task proposed by Howard and Matheson, that of converting an ID into a decision tree, requires the computation of posterior probabilities conditioned on all evidence available at decision time. This computation can be accomplished by any of the standard methods developed for probabilistic inference, e.g., clustering, conditioning, bucket-elimination, or belief propagation.

¹ This interpretation of an arrow in a diagram could not be complete or consistent, for it applies as well to an arrow pointing from B to A , and in fact to any two nodes A and B that are connected via an unblocked path.

² Olmsted (1983) and Shachter (1986) developed procedures for probabilistic inference in the framework of influence diagrams but, lacking conditional independence semantics, were unable to provide performance estimates.

Despite these shortcomings, however, the publication of Howard and Matheson's paper has had a significant impact on decision analytic practices, for it promised to liberate the analyst from the impossible task of mentally estimating the conditional probabilities that decorate the links of decision trees. It also had a catalytic effect on the development of graphical models in general for it intensified the urge to gain formal and conceptual understanding of the then mystical relationships between graphs and probabilities.

Acknowledgments

The research was partially supported by Air Force Grant AFOSR/F49620-03-1-0095, NSF Grant IIS-0097082, and MURI Grant N00014-00-1-0617.

References

- Howard, R. A., J. 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.
- Kelly III, C. W., S. Barklay. 1973. A general Bayesian model for hierarchical inference. *Organ. Behavior Human Performance* 10 388–403.
- Olmsted, S. M. 1983. On representing and solving decision problems. Ph.D. thesis, EES Department, Stanford University, Stanford, CA.
- Pearl, J. 1982. Reverend Bayes on inference engines: A distributed hierarchical approach. *Proc. AAAI National Conf. AI*. Pittsburgh, PA, 133–136.
- Pearl, J. 1986. Fusion, propagation, and structuring in belief networks. *Artificial Intelligence* 29 241–288.
- Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems*. Morgan Kaufmann, San Mateo, CA.
- Pearl, J. 1993. Belief networks revisited. *Artificial Intelligence* 59 49–56.
- Pearl, J. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York.
- Rumelhart, D. E., J. L. McClelland. 1982. An interactive activation model of context effects in letter perception, part 2: The contextual enhancement effect and some tests and extensions of the model. *Psych. Rev.* 89 60–94.
- Shachter, R. D. 1986. Evaluating influence diagrams. *Oper. Res.* 34(6) 871–882.
- Wright, S. 1921. Correlation and causation. *J. Agricultural Res.* 20 557–585.