Introduction to the Special Section on Knowledge-Based Construction of Probabilistic and Decision Models (IEEE Transactions on Systems, Man, and Cybernetics)

John S. Breese
Microsoft Research
Microsoft Corporation
Redmond, WA

Robert P. Goldman
Computer Science Department
Tulane University
New Orleans, LA

Michael P. Wellman
Artificial Intelligence Laboratory
University of Michigan
Ann Arbor, MI

Modeling techniques developed recently in the AI and uncertain reasoning communities permit significantly more flexible specifications of probabilistic knowledge. Specifically, graphical decision-modeling formalisms—belief networks, influence diagrams, and their variants—provide compact representation of probabilistic relationships, and support inference algorithms that automatically exploit the dependence structure in such models [1, 3, 4]. These advances have brought on a resurgence of interest in computational decision systems based on normative theories of belief and preference. However, graphical decision-modeling languages are still quite limited for purposes of *knowledge representation* because, while they can describe the relationships among particular event instances, they cannot capture *general* knowledge about probabilistic relationships across classes of events.

The inability to capture general knowledge is a serious impediment for those AI tasks in which the relevant factors of a decision problem cannot be enumerated in advance. A graphical decision model encodes a particular set of probabilistic dependencies, a predefined set of decision alternatives, and a specific mathematical form for a utility function. Given a properly specified model, there exist relatively efficient algorithms for calculating posterior probabilities and optimal decision policies. A range of similar cases may be handled by parametric variations of the original model. However, if the structure of dependencies, the set of available alternatives, or the form of utility function changes from situation to situation, then a fixed network representation is no longer adequate.

An ideal computational decision system would possess general, broad knowledge of a domain, but would have the ability to reason about the particular circumstances of any given decision problem within the domain. One obvious approach—which we call call *knowledge-based model construction* (KBMC)—is to generate a decision model dynamically at run-time, based on the problem description and information received thus far. Model construction consists of selection, instantiation, and assembly of causal and associational relationships from a broad knowledge base of general relationships among domain concepts.

For example, suppose we wish to develop a system to recommend appropriate actions for maintaining a computer network. The natural graphical decision model would include chance

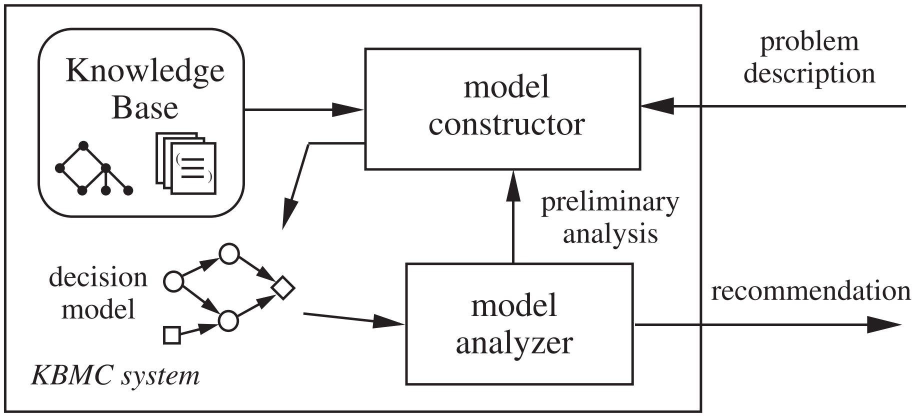

Figure 1: The structure of a knowledge-based model construction system (after [6]). Given a problem instance or query, the system uses general domain knowledge to construct a problem-specific decision model. The constructed model is subsequently solved or used in reasoning tasks.

nodes describing the states of the various networked systems and their components, decision nodes for various test and repair actions, and further chance nodes representing outcomes of these actions. Finally, the model would include a value node representing the overall utility as a function of the state of the system and the various actions taken. The structure of the decision model will depend on the topology of the specific computer network that is being considered, since the observations one can make and the actions one can take depend on the number of servers, clients, and their interconnections.

Now consider the limitations of this model. If one were running a network consulting firm, one would have to compose a new decision model for each client, because each client has a different local network topology. For that matter, one would need a new model for each situation each client encounters, since the relevant actions and events will differ from situation to situation, because various states may cause the network configuration to change.

The KBMC alternative would be to compile a knowledge base containing general knowledge about network function, network maintenance, and the preferences of various clients. For example, we might specify rules indicating that, at some cost, any component of type X can be metered. We would also have rules reflecting how much given clients are willing to pay to avoid interruptions of service. To examine a particular client's network situation, we would provide an input description of the network topology and components. In response the KBMC system would generate a decision model, which on evaluation would compute an optimal policy of maintenance actions.

A typical structure for a KBMC system is shown in Figure 1. The basic approach is to take general, flexible representations of dependency, alternatives, and preferences and combine them with a problem instance or query to develop a decision model tailored to the specific situation. We then use this constructed model as the basis for further inference and decision making.

In the summer of 1991, the American Association for Artificial Intelligence sponsored a workshop with the same title as this special section of the *Transactions*. At this workshop, attendees from AI and the decision sciences identified a set of research issues central to extending techniques for automatically constructing decision models. This special issue presents results from some of the research that grew out of those discussions. Although we emphasize issues revolving about the construction of decision-theoretic models encoded as belief networks or influence diagrams, the issues discussed here also have relevance to construction of simulation models, mathematical programs [2], device models for diagnosis, or any of various other analytical models employed in operations research, management science, artificial intelligence, and related fields.

Core issues in knowledge-based model construction include [6]:

Knowledge Representation What are appropriate fundamental representations of probabilistic relationships, preferences, and alternatives? Most KBMC researchers have used languages based on first-order logic to represent modular components of graphical decision models. Qualitative representations of relative likelihood have also been used, as well as other methods for varying the level of abstraction in a model.

Triggering and Elaboration What triggers the construction process? How does the system expand or contract a model? Current model construction approaches are driven in one of three ways: either by changes in observed world state (data driven), by user query (decision driven), or in response to a specification of user preferences (value driven).

Control of Construction Given a construction procedure, how do we know when it has constructed a model with enough detail (or too much)? Control of construction may require explicit meta-level knowledge about the construction process, in particular an analysis of the benefits and the costs of expanding or contracting a model. Typically there is a tradeoff between the level of detail in a constructed model and the cost of evaluating it.

Knowledge Base Development Given a knowledge representation, how does one specify information in this base language? What guarantees can one make about the quality of the models generated by a knowledge-based construction system? To date, little attention has been given to specification of model construction source languages and their characteristics.

Subsequent research has addressed these and other issues in knowledge-based model construction. The papers included in this special section represent some of the latest advances in this area, including a range of novel techniques and analytical results. In the remainder of this introduction, we summarize the contribution of each paper.

"Graph-grammar assistance for automated generation of influence diagrams," by John W. Egar and Mark A. Musen, provides a novel approach to knowledge representation for KBMC systems. As mentioned above, most existing KBMC systems have used knowledge bases encoded in languages based on first-order logic (often Horn clauses). Egar and Musen propose an alternative representation, based on graph grammars, and demonstrate the use of graph grammars in an interactive program to support medical decision making. Egar and Musen's article also exemplifies an important tactic in KBMC research: to establish analytically that models generated will satisfy desired properties. Egar and Musen show that their influence diagrams are guaranteed not to violate some important criteria for well-designed decision models.

"Inference-driven construction of valuation systems from first-order clauses," by Alessandro Saffioti and Elisabeth Umkehrer, directly addresses the task of constructing valuation systems from knowledge encoded in a logic based on Dempster-Shafer belief functions. Although these are not, strictly speaking, decision-theoretic models, the core KBMC-related issues addressed by this work are isomorphic. Saffiotti and Umkehrer describe a comprehensive model construction system with an architecture similar to that displayed in Figure 1. Their system imposes a clear distinction between the knowledge representation, inference-driven construction, and the solution process. The paper's most important contribution is the formulation and proof of soundness and completeness results for the construction procedure. These theorems get to the heart of what is represented and generated in a model construction system. An interesting area for further research is whether the system described in this paper can be extended to graphical decision models incorporating concepts of probabilistic independence while retaining the same soundness and completeness properties.

Gregory M. Provan's paper, "Tradeoffs in knowledge-based construction of probabilistic models," explores the issue of controlling model construction. Provan suggests several metrics that could be applied to determine whether a decision model has been developed to an appropriate level of detail. He illustrates the use of these metrics to evaluate dynamic decision models for medical diagnosis. Dynamic decision problems pose a particularly difficult challenge for existing problem representations, as naive attempts to capture time-varying processes inevitably lead to models that are impractically large. The challenge is to simplify the dynamic model in ways that sacrifice as little decision quality as possible.

"Dynamic belief networks for discrete monitoring," by A. E. Nicholson and J. M. Brady, presents an interesting case study in dynamic probabilistic modeling. In their robot monitoring domain, sensors generate a stream of observations providing uncertain evidence about the changing location of robot vehicles. Their monitoring system tracks these vehicles via an evolving probabilistic network model, updated and extended at each time step. Their paper considers several strategies to control the size and complexity of the network.

"Metareasoning and the problem of small worlds," by Kathryn Blackmond Laskey and Paul E. Lehner, contributes to an emerging foundational theory of model construction. Savage used the term "small world" to refer to the restricted set of factors that can practically be considered in any formal decision analysis [5]. In their paper, Laskey and Lehner provide a meta-decision-theoretic account of small worlds by explicitly relating the outputs of a small-world model to the output that could be obtained using some larger, more comprehensive world model. They identify techniques for diagnosing when one's model is likely to be faulty and when one should seek an alternate model. These results can inform KBMC by providing ways to control model construction.

"Dynamic construction and refinement of utility-based categorization models," by Kim-Leng Poh, Michael R. Fehling, and Eric J. Horvitz, illustrates the use of hierarchical abstraction in KBMC. The specific task considered in the paper—what the authors call "utility-based categorization"—is to choose a classification label for a situation so as to maximize expected utility. Since the domain of classifications can vary, the choice of a "categorization model" is a form of model construction decision. Poh et al. employ a hierarchy of conceptual covers to manage the flexible choice of abstraction levels. They also show how a meta-level decision model can be used to optimize the construction of a categorization model.

In addition to these works, we are aware of several other ongoing research projects within the KBMC framework. For examples, see recent Proceedings of the *Conference on Uncertainty in Artificial Intelligence*. Although investigations along these lines are just beginning, the preliminary results seem to indicate that automatic generation is a viable approach for extending the scope of graphical decision-modeling technology.

Acknowledgments

We thank the reviewers for this special section for their conscientious evaluations and constructive suggestions to the authors.

References

[1] E. Charniak. Bayesian networks without tears. *AI Magazine*, 12(4):50–63, 1991.

[2] R. Krishnan, Ramayya and X. Li, and D. Steier. A knowledge-based mathematical model formulation system. *Communications of the ACM*, 35(9):138–146, 1992.

[3] R. E. Neapolitan. *Probabilistic Reasoning in Expert Systems: Theory and Algorithms.* John Wiley and Sons, New York, 1990.

[4] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann, San Mateo, CA, 1988.

[5] L. J. Savage. *The Foundations of Statistics.* Dover, New York, 1972. 2nd edition, First edition 1954.

[6] M. P. Wellman, J. S. Breese, and R. P. Goldman. From knowledge bases to decision models. *Knowledge Engineering Review*, 7(1):35–53, 1992.