Building Probabilistic Networks: "Where Do the Numbers Come From?" Guest Editors' Introduction

Marek J. Druzdzel and Linda C. van der Gaag

1 Introduction

Probabilistic networks are now fairly well established as practical representations of knowledge for reasoning under uncertainty, as demonstrated by an increasing number of successful applications in such domains as (medical) diagnosis and prognosis, planning, vision, information retrieval, and natural language processing. A probabilistic network (also referred to as a belief network, Bayesian network, or, somewhat imprecisely, causal network) consists of a graphical structure, encoding a domain's variables and the qualitative relationships between them, and a quantitative part, encoding probabilities over the variables [29].

Building a probabilistic network for a domain of application involves three tasks. The first of these is to identify the variables that are of importance, along with their possible values. Once the important domain variables have been identified, the second task is to identify the relationships between the variables discerned and to express these in a graphical structure. The tasks of eliciting the variables and values of importance, as well as the relationships between them, from domain experts is comparable, to at least some extent, to knowledge engineering for other artificial-intelligence representations and, although it may require significant effort, is generally considered doable. The last task in building a probabilistic network is to obtain the probabilities that are required for its quantitative part. This task often appears more daunting: "Where do the numbers come from?" is a commonly asked question. The three tasks in building a probabilistic network are, in principle, performed one after the other. Building a network, however, often requires a careful trade-off between the desire for a large and rich model to obtain accurate results, on the one hand, and the costs of construction and maintenance and the complexity of probabilistic inference on the other hand. In practice, therefore, building a probabilistic network is a process that iterates over these tasks until a network results that is deemed requisite.

- *M.J. Druzdzel is with the School of Information Science, University of Pittsburgh, Pittsburgh, PA 15260. E-mail: marek@sis.pitt.edu.*
- *L.C. van der Gaag is with the Institute of Information and Computer Sciences, Utrecht University, PO Box 80.089, 3508 TB Utrecht, The Netherlands. E-mail: linda@cs.uu.nl.*

For information on obtaining reprints of this article, please send e-mail to: tkde@computer.org, and reference IEEECS Log Number 112202.

In collaboration with Finn V. Jensen and Max Henrion, we organized in 1995 a workshop devoted to the theme of obtaining the numbers, the most daunting task in building probabilistic networks [14]. The workshop was held in conjunction with the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI '95) and had a program of presentations of selected contributions and ample slots for flash communications and discussion. Scientists from such disciplines as decision analysis, statistics, and computer science attended the workshop. The interest in the workshop, both during IJCAI '95 and afterward, prompted us to follow up on the theme. The current special section of *IEEE Transactions on Knowledge and Data Engineering* is the result.

2 Sources of Probabilistic Information

In most application domains, probabilistic information is available from various sources. The most common are (statistical) data, literature, and human experts. Despite the abundance of information, these sources seldom provide all numbers required for the quantitative part of a probabilistic network. As a consequence, the task of obtaining the numbers for a real-life application is hard and time consuming.

In data-rich application domains, often large data collections are available, retrospectively documenting every-day problem solving. Once the part of the domain to be modeled is well-defined and well-demarcated, it also is not too hard to prospectively collect data on the variables of interest. These data will usually contain highly valuable information about the relationships between the variables in the domain. If a comprehensive data collection is available, the construction of both the graphical part and the quantitative part of a probabilistic network can be performed automatically. The basic idea of the former is to distill information about the relationships between the variables from the data and exploit it for constructing the network's graph. There are essentially two approaches to learning the graphical structure from data. The first is based on constraint-based search [30], [36] and the second on Bayesian search for graphs with highest posterior probability given the data [3]. Once the graphical structure has been established, assessing the required probabilities is quite straightforward and amounts to studying subsets of the data that satisfy various conditions.

To allow for automated construction of a meaningful probabilistic network, the data must have been collected very carefully. Biases that are introduced in the data as a result of the data collection strategies used will usually have an effect on the resulting network [25]. This effect may not be desirable, however, for the purpose for which the network is being developed. Unfortunately, selection biases are not easily detected in a network once it has been constructed. Also, the variables and associated values that are recorded in the data collection should match the variables and values that are to be modeled in the network or should at least admit transformation into these variables and values without too much loss of information [26]. The data collection should further be comprised of enough data to allow for reliable identification of probabilistic relationships among the variables discerned and to provide for reliable probability assessments. In an insufficiently large data collection, the various subsets from which probabilities are estimated, for example, can be empty or too small to allow for meaningful assessments. A common problem typically found in real-life data, especially when it has been retrospectively collected, is the occurrence of missing values. Sometimes a missing value is the result of an error of omission. Quite often, however, a value is not recorded because the variable's measurement did not make sense in practice given the values of other variables. Missing values of the first type are often randomly distributed. Missing values of the second type, on the other hand, generally are not distributed evenly; as a consequence, they are information bearing and need be handled accordingly [31]. To use a data collection with missing values for automated construction of a probabilistic network, often values have to be filled in, for example based upon (roughly) estimated prior or posterior probabilities for these values or with the help of domain experts [6], [16], [34]. Automated construction of probabilistic networks from data is an active area of research [1], [7], [19].

Literature often provides abundant probabilistic information. For every medical diagnostic test, for example, its sensitivity and specificity characteristics, as well as its typical ranges, are reported in medical handbooks or journals. Medical disorders and symptoms, as well as the (causal) relationships between them, are also discussed in ample detail. Unfortunately, the reported probabilistic information is seldom directly amenable to encoding in a probabilistic network. Medical literature, for example, often reports conditional probabilities of the presence of symptoms given a disorder, but not always the probabilities of these symptoms occurring in the absence of the disorder. Also, conditional probabilities are sometimes given in a direction reverse to the direction required for the network. For example, the statement "70 percent of the patients with esophageal cancer are smokers" specifies the probability of a patient being a smoker given that he or she is suffering from esophageal cancer, while, for the network, the probability of esophageal cancer developing in a smoker would be required. Moreover, probabilities for unobservable intermediate disease states are usually lacking altogether. As a consequence, if the reported probabilistic information can be exploited at all, it often requires considerable processing and additional domain knowledge [23]. Another commonly found problem that prohibits direct use of probabilistic information from literature pertains to the characteristics of the population from which the information is derived. These characteristics often are not properly described or deviate seriously from the characteristics of the population for which the probabilistic network is being developed [11]. Almanacs, morbidity and mortality tables, and statistical yearbooks generally suffer less from the problems outlined above. These sources tend to contain fairly reliable probabilistic information that can be used whenever the target population is not atypical.

Finally, when there are few or no reliable data available, the knowledge and experience of experts in the domain of application constitute the only remaining source of probabilistic information. The role of domain experts in the construction of the quantitative part of a probabilistic network should not be underestimated. An expert's knowledge and experience can help, not just in assessing the probabilities required, but also in fine tuning probabilities obtained from other sources to the specifics of the domain at hand and in verifying the numbers within the context of the network. The problems encountered when directly eliciting probabilities from experts, however, are widely known [20]. An expert's assessments, for example, may reflect various biases and may not be properly calibrated. Acknowledging these problems, in the field of decision analysis various techniques have been developed for the elicitation of well-calibrated probabilities from experts, ranging from the use of probability scales for marking assessments to the use of lotteries [27], [43]. These techniques tend to be quite time-consuming and can take up to 30 minutes per number, including the typical overhead in interviews with domain experts, for example, of explaining context. They have found widespread use in the construction of decision-analytic models, which traditionally comprise a reasonably small number of variables. Probabilistic networks tend to differ from conventional decision-analytic models by the number of probabilities they require: Contemporary networks typically comprise tens or hundreds of variables and hundreds or thousands of probabilities. Given that an expert's time is a scarce and expensive commodity, application of the decision-analytic techniques for probability elicitation rapidly becomes impractical if not impossible for network quantification. For probability elicitation for probabilistic networks, therefore, supplementary techniques are being sought [13], [35], [40].

To conclude our brief discussion of sources of probabilistic information, we would like to note that, although tempting, combining information from different sources in a single probabilistic network can be risky and can in fact, lead to incorrect results [10].

3 The Importance of Accurate Numbers

Although generally various sources of information can be exploited for probability assessment, the numbers obtained are inevitably inaccurate due to incompleteness of data and partial knowledge of the domain under study. As the numbers are an integral part of a probabilistic network, their inaccuracies will influence the network's output. It is a natural question, therefore, to ask how accurate the numbers should be to arrive at satisfactory behavior of the network. Experience with constructing probabilistic networks for various domains of application has established a consensus that the graphical structure of a network is its most important part as it reflects the independence and relevance relationships between the variables concerned, being the most robust, qualitative properties of the domain under study [8], [12], [39]. Within the context of the graphical structure, however, numerical inaccuracies will influence the network's output.

The extent to which the inaccuracies in its numbers influence the output of a probabilistic network can be studied by investigating the extent to which deviations from the numbers affect the output. To this end, the network can be subjected to a sensitivity analysis and an uncertainty analysis. In general, sensitivity analysis of a mathematical model amounts to investigating the effects of inaccuracies in the model's parameters on its output by systematically varying the parameters' values [27]. For a probabilistic network, sensitivity analysis amounts to varying the assessments for one or more probabilities in the network's quantitative part simultaneously and investigating the effects on a probability of interest. In an uncertainty analysis, the assessments for all probabilities are varied simultaneously by drawing, for each of them, a value from a prespecified distribution. Uncertainty analysis serves to reveal the overall reliability of a network's output, yet yields less insight into the effect of separate probabilities than a sensitivity analysis.

Uncertainty analysis of a large real-life probabilistic network for liver and biliary disease has provided evidence that probabilistic networks can be highly insensitive to inaccuracies in the numbers in their quantitative part [18], [33]. There is additional, sometimes anecdotal, evidence that networks that contain crude assessments for their probabilities exhibit reasonable behavior. From this evidence, numbers may be looked upon as merely convenient order of magnitude approximations of the strengths of influences between variables. However, evidence is building up that probabilistic networks can be sensitive to the inaccuracies in their numbers. Sensitivity analysis of a real-life network for congenital heart disease, for example, has revealed large effects on a probability of interest [4]. We feel that, from the limited available evidence, no decisive conclusions can be drawn with respect to the effects of inaccuracies in a network's probabilities. At present, it seems likely that these effects will vary from application to application. Sensitivity analysis and uncertainty analysis of probabilistic networks constitute an active field of research that has yielded efficient computational methods for studying the robustness of a network's output [2], [22], [24], [38]. With these methods, more experimental results of sensitivity and uncertainty analyses of real-life probabilistic networks are likely to become available in the near future.

4 Reducing the Burden

Typical contemporary probabilistic networks comprise tens or hundreds of variables, easily requiring thousands of probabilities. It is the vast number of probabilities required that generally hampers the construction of a network for a real-life application. Often the majority of these probabilities have to be assessed by domain experts. As we have argued before, the conventional decision-analytic techniques for probability elicitation are too time-consuming to be suitable for the task. In fact, any contemporary or future technique that aims at eliciting well-calibrated and unbiased probability assessments from domain experts is likely to suffer from this problem. We feel, therefore, that research efforts aimed at reducing the number of probabilities to be assessed and at procedures and tools for supporting the quantification task currently are of more practical significance.

The number of probabilities required for a probabilistic network depends directly on the network's graphical structure. Roughly speaking, the more densely connected a network's graph, the more numbers it requires for its quantitative part. For each variable, exponentially many probabilities have to be provided, their number being exponential in the size of the variable's parental set. There are essentially two approaches to reducing the number of probabilities that have to be assessed for a network. The first is based on changes to the graphical structure and the other on the use of parametric probability distributions. The former approach builds, for example, on the principle of divorcing parents by introducing intermediate variables [28] and on removal of arcs representing weak dependences [21], [41]. The use of a parametric probability distribution for a variable is aimed at reducing the number of probabilities that have to be assessed directly by providing simple rules for the computation of the other probabilities required. Examples of parametric probability distributions currently in use are modeled by the noisy-OR and noisy-AND gates and their generalizations [9], [15], [17], [29], [37]. These models are based on (inter)causal independence assumptions for a variable and its parents. The number of probabilities to be assessed directly for a variable with such a model is linear, rather than exponential, in its number of parents; the remaining, exponentially many, probabilities are readily derived from the independence assumptions underlying the model. A noisy-OR gate, for example, for a binary variable with n binary parents requires n, rather than 2^n, numbers; for $n = 10$, this means a reduction of the number of probabilities to be assessed directly by two orders of magnitude. Changes to the graphical structure of a probabilistic network and the use of parametric distributions are likely to come at the price of accuracy. There currently is little insight into whether or not a fully detailed network with separately specified assessments has a better performance than a network that is carefully reduced using the approaches outlined above. There is no doubt, however, that the reduced network will have required considerably less time on the part of the experts involved. The time thus saved can be exploited for verifying and refining the network.

Building a probabilistic network requires a careful trade-off between the desire for a large and rich model on the one hand and the costs of construction, maintenance, and inference on the other hand. As we have argued before, building a network is a creative and iterative process.

Although obtaining the numbers for a probabilistic network is generally postponed until its graphical structure is considered robust, it is not realistic to assume that the assessment of all numbers required is a one-shot process. Building upon this insight, various research efforts aim at iterative procedures and associated tools to support the daunting quantification task. The procedures build, for example, on the use of sensitivity and value-of-information analyses [5], [32]. As these give insight into the level of accuracy that is required for the various probabilities of a network, they help in focusing elicitation efforts. A procedure building upon sensitivity analysis, for example, sets out with the elicitation of crude, probably highly inaccurate, numbers, within a short period of time [40]. Starting with these numbers, a sensitivity analysis of the network is performed. The most influential probabilities are uncovered, which are thereupon refined, for example, using conventional decision-analytic elicitation techniques. As a side-effect, the analysis can point to uninfluential parts of a network that may be deleted or simplified. Iteratively performing sensitivity analyses and refining probabilities is pursued until satisfactory behavior of the network is obtained, until the costs of further elicitation outweigh the benefits of higher accuracy, or until higher accuracy can no longer be attained due to lack of knowledge. Given the limited and costly time of experts, attention can thus be focused on the probabilities to which the network's behavior shows highest sensitivity.

Procedures for network quantification can be supported by graphical tools that provide for interactive elicitation, inspection, and modification of probabilities [42]. With such tools, probabilities can be elicited through a variety of modalities. Direct elicitation of probabilities, while easiest to implement, is generally the least reliable. Elicitation using graphical means that allow an expert to directly manipulate a pie chart or a bar graph offers more support to the expert and is likely to lead to numbers with higher accuracy. Probabilities can also be related to verbal descriptions such as *very likely* and *improbable* [35]. Although verbal descriptions of probabilities are known to be context sensitive and can describe wide ranges of numerical quantities, the use of both words and numbers in probability elicitation can result in reasonable assessments [40]. Rather than pushing an expert to assess a large number of probabilities, tools for interactive probability elicitation can support noninvasive elicitation by accommodating whatever probabilistic information the expert is willing to provide [13]. This information may be quantitative in nature, such as point estimates and probability intervals, but may also be qualitative, such as comparisons and statements of stochastic dominance. Procedures for network quantification can furthermore be supported by tools for automated generation of explanations of reasoning behavior. Detailed explanation provides for studying the reasoning behavior of a network, which can point to problems with the probabilities in the network's quantitative part.

To conclude, in this brief introduction, we have focused attention on the task of obtaining the numbers required for a probabilistic network as this is the main scope of the current issue. However, as we have argued before, the quantification task is not performed in isolation from the rest of the process of building a network. With the increasing number of applications, a need for knowledge-engineering principles tailored to the construction of probabilistic networks is emerging. With the advance of iterative procedures and associated graphical tools for supporting the overall construction process, the quantification task will be addressed within its proper context, which hopefully will contribute to reducing its burden.

5 Guide to the Contents

Limited space allows for a special section of a journal to focus only on selected aspects of a problem. The current issue of *IEEE Transactions on Knowledge and Data Engineering* is no exception. The papers that we have selected cover just some of the topics addressed above; each of the papers, however, addresses one or more of these topics in detail and does so from the point of view of experiences with building real-life probabilistic networks. As a guide to the contents of the section, we briefly review the four selected papers.

In "Network Engineering for Agile Belief Network Models," K. Blackmond Laskey and S.M. Mahoney argue that the quantification task is best treated within the broader context of building probabilistic networks because of the interplay between structural modeling decisions and the numbers to be obtained. The authors propose an integrated systems engineering approach in building probabilistic networks to take this interplay into consideration. In "Dealing with the Expert Inconsistency in Probability Elicitation," S. Monti and G. Carenini describe experiences with elicitation of probabilities from human domain experts for a real-life probabilistic network in the domain of chronic nonorganic headaches. The authors evaluate and compare the use of various techniques for probability elicitation, among which is a newly developed technique that provides for uncovering inconsistencies in an expert's assessments. D. Nikovski in "Constructing Bayesian Networks for Medical Diagnosis from Incomplete and Partially Correct Statistics" introduces several knowledge engineering techniques for the construction of probabilistic networks, tailored to the domain of medical diagnosis. The author focuses on the task of obtaining the numbers required for a network from available probabilistic information that is incomplete. In the last paper of the section, "A Causal Probabilistic Network for Optimal Treatment of Bacterial Infections," L. Leibovici, M. Fishman, H.C. Schønheyder, C. Riekehr, B. Kristensen, I. Shraga, and S. Andreassen describe their experiences with building a large real-life probabilistic network for the treatment of severe bacterial infections. The probabilities in the network's quantitative part have been obtained from large data collections and from literature. The authors address the structural decisions they have made to take account of the availability of data from which the required probabilities could be estimated.

Acknowledgments

We would like to thank the participants of the IJCAI 1995 Workshop on Building Probabilistic Networks: Where Do the Numbers Come from? for their active participation and stimulating discussions. We would also like to thank the numerous reviewers for their generous contribution to the resulting quality of this special section. Dr. Druzdzel was supported by the US National Science Foundation under the Faculty Early Career Development (CAREER) Program, grant IRI-9624629, and by the US Air Force Office of Scientific Research, grants F49620-97-1-0225 and F49620-00-1-0112.

References

[1] W. Buntine, "A Guide to the Literature on Learning Graphical Models," *IEEE Trans. Knowledge and Data Eng.*, vol. 8, pp. 195-210, 1996.

[2] E. Castillo, J.M. Gutiérrez, and A.S. Hadi, "Sensitivity Analysis in Discrete Bayesian Networks," *IEEE Trans. Systems, Man, and Cybernetics,* vol. 27, pp. 412-423, 1997.

[3] G.F. Cooper and E. Herskovits, "A Bayesian Method for the Induction of Probabilistic Networks from Data," *Machine Learning,* vol. 9, pp. 309-347, 1992.

[4] V.M.H. Coupé, N. Peek, J. Ottenkamp, and J.D.F. Habbema, "Using Sensitivity Analysis for Efficient Quantification of a Belief-Network," *Artificial Intelligence in Medicine,* vol. 17, pp. 223-247, 1999.

[5] V.M.H. Coupé, L.C. van der Gaag, and J.D.F. Habbema, "Sensitivity Analysis: An Aid for Probability Elicitation," *Knowledge Eng. Review,* to appear.

[6] R.G. Cowell, "Parameter Estimation from Incomplete Data for Bayesian Networks," *Proc. Seventh Int'l Workshop Artificial Intelligence and Statistics,* pp. 193-196, 1999.

[7] R.G. Cowell, A.P. Dawid, S.L. Lauritzen, and D.J. Spiegelhalter, *Probabilistic Networks and Expert Systems.* Berlin-Heidelberg-New York: Springer-Verlag, 1999.

[8] A.P. Dawid, "Conditional Independence in Statistical Theory," *J. Royal Statistical Soc., Series B (Methodological),* vol. 41, pp. 1-31, 1979.

[9] F.J. Díez, "Parameter Adjustment in Bayes Networks: The Generalized Noisy OR-Gate," *Proc. Ninth Conf. Uncertainty in Artificial Intelligence,* pp. 99-105, 1993.

[10] M.J. Druzdzel and F.J. Díez, "Criteria for Combining Knowledge from Different Sources in Probabilistic Models," *Working Notes of the UAI 2000 Workshop on Domain Knowledge with Data for Decision Support,* to appear.

[11] M.J. Druzdzel, A. Onisko, D. Schwartz, J.N. Dowling, and H. Wasyluk, "Knowledge Engineering for Very Large Decision-Analytic Medical Models," *Proc. 1999 Ann. Meeting Am. Medical Informatics Assoc. (AMIA-99),* p. 1,049, 1999.

[12] M.J. Druzdzel and H.J. Suermondt, "Relevance in Probabilistic Models: 'Backyards' in a Small World," *Working Notes AAAI-1994 Fall Symp. Series: Relevance,* pp. 60-63, 1994.

[13] M.J. Druzdzel and L.C. van der Gaag, "Elicitation of Probabilities for Belief Networks: Combining Qualitative and Quantitative Information," *Proc. 11th Conf. Uncertainty in Artificial Intelligence,* pp. 141-148, 1995.

[14] *Working Notes IJCAI '95 Workshop Building Probabilistic Networks: Where Do the Numbers Come From?* M.J. Druzdzel, L.C. van der Gaag, M. Henrion, and F.V. Jensen, eds. Menlo Park, Calif.: Am. Assoc. Artificial Intelligence, 1995.

[15] D. Heckerman and J.S. Breese, "Causal Independence for Probabilistic Assessment and Inference Using Bayesian Networks," *IEEE Trans. Systems, Man, and Cybernetics,* vol. 26, pp. 826-831, 1996.

[16] D. Heckerman, D. Geiger, and D.M. Chickering, "Learning Bayesian Networks: The Combination of Knowledge and Statistical Data," *Machine Learning,* vol. 20, pp. 197-243, 1995.

[17] M. Henrion, "Some Practical Issues in Constructing Belief Networks," *Uncertainty in Artificial Intelligence 3,* L.N. Kanal, T.S. Levitt, and J.F. Lemmer, eds., pp. 161-173, North Holland: Elsevier Science Publishers B.V., 1989.

[18] M. Henrion, M. Pradhan, B. Del Favero, K. Huang, G. Provan, and P. O'Rorke, "Why Is Diagnosis Using Belief Networks Insensitive to Imprecision in Probabilities?" *Proc. 12th Conf. Uncertainty in Artificial Intelligence,* pp. 307-314, 1996.

[19] M.I. Jordan, *Learning in Graphical Models.* Dordrecht: Kluwer Academic, 1998.

[20] D. Kahneman, P. Slovic, and A. Tversky, *Judgment under Uncertainty: Heuristics and Biases.* Cambridge, U.K.: Cambridge Univ. Press, 1982.

[21] U. Kjærulff, "Reduction of Computational Complexity in Bayesian Networks through Removal of Weak Dependences," *Proc. 10th Conf. Uncertainty in Artificial Intelligence,* pp. 374-382, 1994.

[22] U. Kjærulff and L.C. van der Gaag, "Making Sensitivity Analysis Computationally Efficient," *Proc. 16th Conf. Uncertainty in Artificial Intelligence,* to appear.

[23] M. Korver and P.J.F. Lucas, "Converting a Rule-Based Expert System into a Belief Network," *Medical Informatics,* vol. 18, pp. 219-241, 1993.

[24] K.B. Laskey, "Sensitivity Analysis for Probability Assessments in Bayesian Networks," *IEEE Trans. Systems, Man, and Cybernetics,* vol. 25, pp. 901-909, 1995.

[25] P.J.F. Lucas, H. Boot, and B.G. Taal, "Computer-Based Decision-Support in the Management of Primary Gastric Non-Hodgkin Lymphoma," *Methods of Information in Medicine,* vol. 37, pp. 206-219, 1998.

[26] P.J.F. Lucas, N.C. de Bruijn, K. Schurink, and I.M. Hoepelman, "A Probabilistic and Decision-Theoretic Approach to the Management of Infectious Disease at the ICU," *Artificial Intelligence in Medicine,* vol. 19, pp. 251-279, 2000.

[27] M.G. Morgan and M. Henrion, *Uncertainty, a Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis.* Cambridge, U.K.: Cambridge Univ. Press, 1990.

[28] K.G. Olesen, U. Kjærulff, F. Jensen, F.V. Jensen, B. Falck, S. Andreassen, and S.K. Andersen, "A MUNIN Network for the Median Nerve—A Case Study on Loops," *Applied Artificial Intelligence,* vol. 3, pp. 385-404, 1989.

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

[30] J. Pearl and T.S. Verma, "A Theory of Inferred Causation," *Principles of Knowledge Representation and Reasoning: Proc. Second Int'l Conf.,* J.A. Allen, R. Fikes, and E. Sandewall, eds., pp. 441-452, 1991.

[31] M.A. Peot and R.D. Shachter, "Learning from What You Don't Observe," *Proc. 14th Conf. Uncertainty in Artificial Intelligence,* pp. 439-446, 1998.

[32] L.D. Philips, "Requisite Decision Modeling: A Case Study," *J. Operational Research Soc.,* vol. 3, pp. 303-311, 1982.

[33] M. Pradhan, M. Henrion, G. Provan, B. del Favero, and K. Huang, "The Sensitivity of Belief Networks to Imprecise Probabilities: An Experimental Investigation," *Artificial Intelligence,* vol. 85, pp. 363-397, 1996.

[34] M. Ramoni and P. Sebastiani, "Learning Conditional Probabilities from Incomplete Data: An Experimental Comparison," *Proc. Seventh Int'l Workshop Artificial Intelligence and Statistics,* 1999.

[35] S. Renooij and C.L.M. Witteman, "Talking Probabilities: Communicating Probabilistic Information with Words and Numbers," *Int'l J. Approximate Reasoning,* vol. 22, pp. 169-194, 1999.

[36] P. Spirtes, C. Glymour, and R. Scheines, *Causation, Prediction, and Search.* New York: Springer Verlag, 1993.

[37] S. Srinivas, "A Generalization of the Noisy-OR Model," *Proc. Ninth Conf. Uncertainty in Artificial Intelligence,* pp. 208-215, 1993.

[38] L.C. van der Gaag and V.M.H. Coupé, "Sensitivity Analysis for Threshold Decision Making with Bayesian Belief Networks," *AI*IA 99: Advances in Artificial Intelligence,* E. Lamma and P. Mello, eds., pp. 37-48, 2000.

[39] L.C. van der Gaag and J.-J.C. Meyer, "Informational Independence: Models and Normal Forms," *Int'l J. Intelligent Systems,* vol. 13, pp. 83-109, 1998.

[40] L.C. van der Gaag, S. Renooij, C.L.M. Witteman, B.M.P. Aleman, and B.G. Taal, "How to Elicit Many Probabilities," *Proc. 15th Conf. Uncertainty in Artificial Intelligence,* pp. 647-654, 1999.

[41] R.A. van Engelen, "Approximating Bayesian Belief Networks by Arc Removal," *IEEE Trans. Pattern Analysis and Machine Intelligence,* vol. 19, pp. 916-920, 1997.

[42] H. Wang and M.J. Druzdzel, "User Interface Tools for Navigation in Conditional Probability Tables and Elicitation of Probabilities in Bayesian Networks," *Proc. 16th Conf. Uncertainty in Artificial Intelligence,* to appear.

[43] D. von Winterfeldt and W. Edwards, *Decision Analysis and Behavioral Research.* New York: Cambridge Univ. Press, 1986.

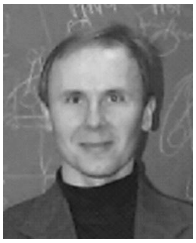

Marek J. Druzdzel received his MS degrees in computer science (1985) and electrical engineering (1987) from Delft University of Technology in The Netherlands (both with distinction) and his PhD in engineering and public policy (1992) from Carnegie Mellon University. He is currently an associate professor in the School of Information Sciences and Intelligent Systems Program at the University of Pittsburgh. He heads the Decision Systems Laboratory, an interdisciplinary research group focusing on normative decision support systems. His research interests concentrate on probabilistic and decision theoretic reasoning in decision support systems and user interfaces to decision support systems. More details about his research interests and publications are available on line at http://www.pitt.edu/~druzdzel.

Linda C. van der Gaag received her MSc degree in mathematics (1983) from Delft University of Technology and her PhD degree in computer science (1990) from the University of Amsterdam in The Netherlands. In 1989, she joined the Institute of Information and Computer Sciences at Utrecht University in the Netherlands where she is currently a full professor, heading the Utrecht research group in decision-support systems. Her research interests concentrate on the principles of reasoning with uncertainty and decision making under uncertainty and on the exploitation of these principles in decision-support systems.