AI Magazine Volume 15 Number 4 (1994) (© AAAI)

Research Issues in Qualitative and Abstract Probability

Moises Goldszmidt

■ To assess the state of the art and identify issues requiring further investigation, a workshop on qualitative and abstract probability was held during the third week of November 1993. This workshop brought together a mix of active researchers from academia, industry, and government interested in the practical and theoretical impact of these abstractions on techniques, methods, and tools for solving complex AI tasks. The result was a set of specific recommendations on the most promising and important avenues for future research.

On 19–20 November 1993, a 2-day workshop on qualitative and abstract probabilities (QAP) was held in San Francisco, California. The workshop, entitled "Putting Qualitative and Abstract Probability to Work," gathered active researchers from university, industry, and government to assess the state of the art and make recommendations for future research. The event was sponsored by the Palo Alto Laboratory of Rockwell Science Center.

Although most of the different QAP formalisms were recently developed, a substantial literature on them already exists. Situated at the crossroad between symbolic knowledge in AI and numeric probabilistic reasoning, QAP formalisms have established a bridge between different areas in knowledge representation, including nonmonotonic and default reasoning; belief change; decision making; planning; and, of course, uncertainty in AI. This cross-fertilization includes notions of conditioning, independence, new algorithms for inference, the use of common benchmarks and examples, and the interchange between research methodologies.

QAPs bring the immediate benefit of allowing the effective representation of uncertain information even when a complete probabilistic distribution (in the classical numeric sense) is unavailable. They are also capable of reasoning with this incomplete information in a fashion consistent with the norms of probability theory. However, a number of important questions need to be addressed: In what applications and instances can QAP formalisms replace standard numeric probabilities or a logical-functional specification? What are the computational benefits of approximating a numeric probability distribution using any of these formalisms? What are the trade-offs between loss of information and computational and representational benefits? In addition, a number of research topics present themselves as promising and important, including the study of frameworks allowing mixtures of qualitative and quantitative information, the characterization of new algorithms that take advantage of the new representational power, the application of these formalisms to tasks such as planning[1] and diagnosis and repair, and the development of compilation schemes.[2]

A report on the 1993 San Francisco workshop

Participants in the workshop included not only researchers in the field of QAPs but also researchers working in applications of numeric probabilities and decision making, such as diagnosis, repair, and planning; logicians; and traditional symbolic AI scientists. The objective was to allow a challenging interchange of ideas among researchers in the field yet also incorporate the concerns, criticisms, and vision of the potential users of these formalisms. To focus the discussion, the workshop centered on three panels: (1) Impact on Uncertainty in AI Methods and Techniques, (2) Impact on Logical Reasoning, and (3) Conditioning: Worth the Price?

The outcome of these panels was a set of well-defined questions and issues requiring further investigation. The workshop concluded with an open-ended session on the expected effectiveness of various approaches and directions in future research.

Qualitative and Abstract Probabilities

Many proposed formalisms can qualify as either a qualitative version or an abstraction of probability theory: qualitative probability networks (Druzdzel and Henrion 1993; Wellman 1990b), infinitesimal probabilities (or epsilon semantics) (Goldszmidt, Morris, and Pearl 1993; Geffner 1992; Pearl 1992) and rankings (Goldszmidt and Pearl 1992a, 1992b), random worlds (Bacchus et. al. 1992; Grove, Halpern, and Koller 1992), argument networks (Darwiche 1993), interval-valued probability distributions (Fertig and Breese 1993), and so on. Some of these formalisms have been used (and even developed) to provide a semantic (and sometimes computational) account for default and nonmonotonic reasoning and formal models for belief change (Pearl 1991). The main concern in this case is the formalization of a notion of plain belief and the ability to produce predictions and explanations in the presence of incomplete information.

Other QAP formalisms are concerned with the process of decision making (Pearl, 1993; Wellman 1990a). Typically, in critical decision-making situations, the information needed to build a complete probabilistic and utility model is unavailable. What is often available instead is either a crude estimate or a vague sense in which certain events increase or decrease our belief in other events. These different motivations and concerns provide a dimension under which we can start building a taxonomy of QAP formalisms, based on, for example, the task the formalisms are trying to accomplish: belief acceptance or decision making. Such an attempt at building a taxonomy for QAPs is investigated by Wellman (1993).

Another dimension for comparison can be provided by the analysis of the formal properties of the QAPs with respect to the two major notions in reasoning systems based on probabilities: a state of belief and conditioning. The axiomatic characterization of these notions proposed by Darwiche and Ginsberg (1992) constitutes an excellent starting point and might lead to a formal definition of a QAP formalism.

In spite of the technical differences, all these formalisms have in common two characteristics: (1) a representation framework and reasoning process consistent with probability theory and (2) the ability to allow belief acceptance and decision making to commence without requiring a single number for each probability and utility. Rather than focus on the issue of a formal definition of what constitutes a QAP formalism, the three panels sparked a set of interesting discussions that resulted in recommendations for future research directed toward increasing the range of applicability of the QAPs to problems in AI.

Panels

The contents and focus of the three panels ranged from practical concerns, including the proposal of a benchmark example for a diagnosis and repair application, to philosophical and theoretical considerations such as the inevitability of conditioning in all these formalisms and its consequences.

The objective of the panel on the impact of QAP on uncertainty in AI methods and techniques was to discuss and assess the role of QAPs in applications requiring reasoning under uncertainty and decision making. The discussion centered on the possible computational and representational advantages of QAPs, the loss of information in regard to belief assessment and its impact on the quality of decision making, and the possibilities of intermixing qualitative and quantitative approaches to engineer systems with a variable degree of specificity[3] or granularity.

A car diagnosis and repair example was proposed as a benchmark by Jack Breese and David Heckerman. This example is fairly complex and includes repair costs, observations, and uncertainties regarding both observations and faults.[4] The objective is to find optimal recommendations for repair. The issues involved in accomplishing this task involve the effective representation of time, persistence, and the interplay between observations and repairs. Numerous assumptions had to be made to make the implementation practical under a standard numeric representation based on Bayes networks (Heckerman, Breese, and Rommelse 1994). Among the questions raised were whether a QAP formalism can relax some of these assumptions and whether the representational characteristics can provide some computational speed and robustness to changes in the degree of belief used to encode the causal relations.

One feature distinguishing QAI's from other KR schemes is their inclusion of a fundamental notion of conditioning.

As mentioned previously, research in QAP initiated a set of formal interactions among different subareas in knowledge representation, including default, uncertainty, and more traditional logical reasoning. The purpose of the panel on the impact of QAP on logical reasoning was to explore the influences and consequences of exporting notions of conditional independence and causal networks to these camps in knowledge representation. These notions are, in general, responsible for the strong presence of probabilistic reasoning in real-world applications. This panel included opposite positions that, on the one hand, challenged the embedded Markovian assumptions in a Bayes network and, on the other hand, assumed them as inevitable for practical systems. There was also a discussion on the theoretical advantages of these assumptions in representing causal relations.

One feature distinguishing QAPs from other knowledge representation schemes is their inclusion of a fundamental notion of conditioning; the panel entitled Conditioning: Worth the Price? focused on the advantages and disadvantages that such a notion provides. Questions raised by the panelists included, What are the costs of representing knowledge using methods supporting conditioning? What are the associated savings in using the knowledge? These questions were addressed from both the standpoint of theoretical results and the standpoint of lessons learned from implemented systems working on practical problems.

Summary

Several of the issues raised in the panels, including questions of different trade-offs between information loss and optimal decision making, average case complexity, and the exploration of QAP formalisms used as heuristic approximations of well-formed probability distributions, pointed toward the need for experimental evaluations and accessible implementations of QAP proposals. To facilitate the interchange of experimental data, implementations, and new results, an anonymous ftp site was created at rpal.rockwell.com under the directory /home2/ftp/pub/QAP.[5] An e-mail list, qap@rpal.rockwell.com, was created as well to encourage further discussions.[6]

The discussions produced a list of suggestions for future research (although some of the issues were already being tackled by a number of participants in the workshop), including the following:

First is the exploration of mixtures of various QAP frameworks with numeric probabilities. There are (at least) two dimensions in which these mixtures should be explored: The first is using QAPs as approximations of available numeric probability distributions. The idea is to have hierarchical abstractions with the objective of speeding up computation (or even getting anytime behavior) in real-time applications. The second is to use QAPs in those cases where the probabilities are unavailable or difficult to obtain (for example, in the assessment of the impact of certain technology or in the likelihood of a surprise attack) in conjunction with numeric probabilities (for example, weather predictions or disease—symptom relations) to produce sensible and coherent answers to user queries.

Second is the exploration of the different intermediate levels of abstraction between a numeric specification of probabilities and any of the QAP specification of uncertainty.

Third are the experimental and theoretical characterizations of the trade-offs involved in abstracting numeric probabilities with respect to loss of information and robustness both in belief assessment and decision making.

Fourth is the identification and characterization of classes of probability distributions or network configurations that exhibit good computational properties, worst case and average case, with regard to inference for specific QAP formalisms. This item includes the study and formulation of specialized algorithms for inference that take advantage of specific representation properties of a particular QAP.

Fifth is the study of the use of QAPs for creating and indexing compiled if-then rules for reactive planning (including diagnosis and repair) and learning. The idea is to take advantage of the robustness of these abstractions with respect to variations in the numeric uncertainty when encoding the domain.

Sixth is the development and extension of current QAP approaches to include notions of time and complex representations of actions, with the objective of incorporating computationally effective notions of rational decision making into AI applications.

In conclusion, the workshop accomplished its objectives and provided participants with the unique opportunity to shape and focus future research in a crucial area for AI applications such as effective knowledge representation and inference methods in decision making and reasoning under uncertainty.

Acknowledgments

Many thanks to Fahiem Bacchus, Jack Breese, Craig Boutilier, Tom Chavez, Adnan Darwiche, Thomas Dean, Marek Druzdzel, Nir Friedman, Matthew Ginsberg, Adam Grove, Joe Halpern, Othar Hansson, David Heckerman, Max Henrion, Eric Horvitz, Daphne Koller, John Lemmer, Judea Pearl, David Poole, Olivier Raiman, David E. Smith, Mike Wellman, and Jim White for their enthusiastic participation in the workshop. Special thanks go to Max Henrion for the idea, Adnan Darwiche for his invaluable help with structuring the panels, and Cathy Froio and Diane Tabejilla for their help with the organization.

Adnan Darwiche, Matthew Ginsberg, and Jim Martin provided comments to previous versions of this article. Peter Patel-Schneider provided useful editorial comments.

Notes

1. See, for example, the work by Wellman (1990a).

2. By compilation schemes, I mean the use of QAPs to create default if-then rules for a reactive planner or suitable tables in learning applications.

3. This term is owed to Max Henrion.

4. The complete set of data and a postscript file with illustrative slides can be retrieved by anonymous ftp from rpal.rockwell.com at /home2/ftp/pub/QAP.

5. Researchers have already started to take advantage of this setup to compare and share experimental findings dealing with applying QAP to diagnosis. Initial results are reported in Darwiche and Goldszmidt (1994) and Henrion et. al. (1994).

6. Requests for additions to the list should be sent to qap-requests@rpal.rockwell.com.

References

Bacchus, F.; Grove, A.; Halpern, J.; and Koller, D. 1992. Statistical Foundations of Default Reasoning Indifference and Irrelevance. Presented at the Fourth International Workshop on Nonmonotonic Reasoning, 28–31 August, Plymouth, Vermont.

Darwiche, A. 1993. Argument Calculus and Networks. In *Proceedings of the Ninth Conference on Uncertainty in Artificial Intelligence,* 420–421. San Francisco, Calif.: Morgan Kaufmann.

Darwiche, A., and Goldszmidt, M. 1994. On the Relation between Kappa Calculus and Probabilistic Reasoning. In *Proceedings of the Tenth Conference on Uncertainty in Artificial Intelligence,* 145–153. San Francisco, Calif.: Morgan Kaufmann.

Darwiche, A., and Ginsberg, M. 1992. A Symbolic Generalization of Probability Theory. In Proceedings of the Ninth National Conference on Artificial Intelligence, 622–627. Menlo Park, Calif.: American Association for Artificial Intelligence.

Druzdzel, M., and Henrion, M. 1993. Efficient Reasoning in Qualitative Probabilistic Networks. In Proceedings of the Tenth National Conference on Artificial Intelligence, 548–553. Menlo Park, Calif.: American Association for Artificial Intelligence.

Fertig, K., and Breese, J. 1993. Probability Intervals over Influence Diagrams. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 15(3): 280–286.

Geffner, H. 1992. *Default Reasoning: Causal and Conditional Theories*. Cambridge, Mass.: MIT Press.

Goldszmidt, M.; Morris, P.; and Pearl, J. 1993. A Maximum Entropy Approach to Nonmonotonic Reasoning. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 15(3): 220–232.

Goldszmidt, M., and Pearl, J. 1992a. Rank-Based Systems: A Simple Approach to Belief Revision, Belief Update, and Reasoning about Evidence and Actions. In *Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference,* 661–672. San Francisco, Calif.: Morgan Kaufmann.

Goldszmidt, M., and Pearl, J. 1992b. Reasoning with Qualitative Probabilities Can Be Tractable. In *Proceedings of the Eighth Conference on Uncertainty in AI,* 112–120. San Francisco, Calif.: Morgan Kaufmann.

Grove, A.; Halpern, J; and Koller, D. 1992. Random Worlds and Maximum Entropy. In Proceedings of the Seventh IEEE Symposium of Logic in Computer Science, 22–23. Washington, D.C.: IEEE Computer Society.

Heckerman, D.; Breese, J.; and Rommelse, K. 1995. Sequential Troubleshooting under Uncertainty. *Communications of the ACM* (Special Issue on Real-World Applications of Uncertain Reasoning). Forthcoming.

Henrion, M.; Provan, G.; del Favero, B.; and Sanders, G. 1994. An Experimental Comparison of Diagnostic Performance Using Infinitesimal and Numerical Bayesian Belief Networks. In *Proceedings of the Tenth Conference on Uncertainty in Artificial Intelligence,* 319–326. San Francisco, Calif.: Morgan Kaufmann.

Pearl, J. 1993. From Conditional Oughts to Qualitative Decision Theory. In *Proceedings of the Ninth Conference on Uncertainty in AI,* 12–20. San Francisco, Calif.: Morgan Kaufmann.

Pearl, J. 1992. Epsilon-Semantics. In *Encyclopedia of Artificial Intelligence,* 468–475. New York: Wiley.

Pearl, J. 1991. Probabilistic Semantics for Nonmonotonic Reasoning: A Survey. In *Philosophy and AI—Essays at the Interface*, eds. R. Cummins and J. Pollock, 57–187. Cambridge, Mass.: Bradford Books.

Wellman, M. 1993. Some Varieties of Qualitative Probabilities, Technical Report, Dept. of Electrical Engineering and Computer Science, Univ. of Michigan.

Wellman, M. 1990a. *Formulation of Trade-Offs in Planning under Uncertainty.* London: Pitman.

Wellman, M. 1990b. Fundamental Concepts of Qualitative Probabilistic Networks. *Artificial Intelligence* 44:257–303.

Moises Goldszmidt is a research scientist with the Information and Decision Sciences Group at the Rockwell Science Center in Palo Alto, California. He received his Ph.D. in computer science from the University of California at Los Angeles in 1992 and was the recipient of an IBM graduate fellowship in 1990 and 1991. His current research interests include reasoning and decision making under uncertainty, tractable and computationally effective methods for knowledge representation and reasoning, diagnosis, and planning.