
 WORKSHOP REPORT

Uncertainty in Artificial Intelligence

Tod S. Levitt

The Fourth Uncertainty in Artificial Intelligence workshop was held 19–21 August 1988. The workshop featured significant developments in application of theories of representation and reasoning under uncertainty. A recurring idea at the workshop was the need to examine uncertainty calculi in the context of choosing representation, inference, and control methodologies.

The effectiveness of these choices in AI systems tends to be best considered in terms of specific problem areas. These areas include automated planning, temporal reasoning, computer vision, medical diagnosis, fault detection, text analysis, distributed systems, and behavior of nonlinear systems. Influence diagrams are emerging as a unifying representation, enabling tool development. Interest and results in uncertainty in AI are growing beyond the capacity of a workshop format.

The Fourth Workshop on Uncertainty in Artificial Intelligence was held 19–21 August 1988 at the University of Minnesota. The workshop featured significant developments in the application of theories of representation and reasoning under uncertainty to diverse areas. These areas included automated planning, temporal reasoning (Dean and Kanazawa, p. 73; Dutta, p. 90), computer vision (Agosta, p. 1; Levitt, Binford, et al., p. 245), medical diagnosis (Cecile et al., p. 38), fault detection (Gallant, p. 127), text analysis (Tessem and Ersland, p. 344), distributed systems (Frisch and Haddawy, p. 109), and behavior of nonlinear systems (Yeh, p. 374).

A central thrust that uncertainty in artificial intelligence (UAI) brings to these applications is the pressing need to handle real-world combinatorics. Here, UAI techniques can make a unique contribution to the AI world. For example, Kyburg (p. 229) suggested that nonmonotonic reasoning might be subsumed by appropriately formulated probabilistic inference, bypassing the combinatorics inherent in purely logical representations; Yager (p. 368) presented a similar scheme for possibility theory. Neufeld and Poole (p. 275) suggested that an ordinal, qualitative calculus, might be sufficient to deal with nonmonotonic and default representations. However, the theoretical results of Aleliunas (p. 8) were evidence that calculi with any ordinal value space might be forced to pure probability. Kadie (p. 197), Bachus (p. 15), and Frisch and Haddawy (p. 109) also presented methods for integrating probabilistic and logical reasoning, both to limit combinatorics and maximize the expressiveness of representations. In the same

spirit, Fu (p. 119) presented a method of nonmonotonic reasoning in a neural net representation.

A theme from the 1987 workshop, control of inference processes, is developed both as a method of limiting combinatorics and as an inherent component of AI systems. D'Ambrosio (p. 64) and Jimison (p. 189) presented methods for modular, interactive system modeling, inference, and control. Utility theory provides a control approach for systems whose uncertain inference is represented in a Bayesian framework. Levitt, Binford, et al. (p. 245), Breese and Fehling (p. 30), Hanson and Mayer (p. 148), and Horvitz (unpublished) provided implementation techniques and efficiency arguments for utility-based control. Pittarelli (p. 283) and Wen (p. 360) presented similar arguments about minimum entropy approaches. Paul Cohen (unpublished) gave a persuasive presentation for the integration of procedural (for example, reactive) representations, with numeric calculi, in the context of an automated planner for fire fighting control. Jain and Agogino (p. 178) and Eick (p. 98) presented fuzzy approaches to the representation and control of systems. Here, the emphasis was as much on issues of the sensitivity of variables in presentations and their influence on control of systems as it was on the dichotomy between fuzzy and probabilistic representations.

The comparison and fusion of multiple UAI theories continued to be central issues. Neapolitan and Kenevan (p. 266) compared classical to Bayesian probability; Kwok and Carter (p. 213) compared averaging effects in decision trees, and Kalagnanam and Henrion (p. 205) compared Bayesian to heuristic tech-

niques. Black and Laskey (p. 22), Heckerman (p. 158), and Loui (p. 257) compared various probabilistic methods to each other and Dempster-Shafer (D-S) approaches. Shenoy and Shafer (p. 307) discussed their axiomatic framework that incorporates belief and probability frameworks. Dubois and Prade (p. 81) and Yen (p. 382) presented fusion models for fuzzy and D-S approaches.

A recurring idea at the workshop was the need to examine calculi in the context of choices of representation, inference, and control methodologies. The effectiveness of these choices in AI systems tends to be best considered in terms of specific problem areas. For this reason, it is difficult to make meaningful, yet general, statements about UAI. The choice of problem structuring and representation using influence diagrams is obvious for works dealing directly with belief revision in probabilistic inference, such as those of Cooper (p. 55); Reid, Parnell, and Morlan (p. 291); Suermondt and Cooper (p. 335); and

Shachter (p. 299). Shachter presented some new techniques for the efficient performance of probabilistic inference. However, it is notable that Jain and Agogino and Shenoy and Shafer, who are engaged in work on possibility and belief function approaches, respectively, presented their work using influence diagram representations. Influence diagrams were also used in Hunter's (p. 170) work on Spohn's (p. 315) nonprobabilistic calculi and in other presentations. It seems that influence diagrams, traditionally associated with probabilistic and decision-theoretic approaches, have become accepted as a general and effective representation and modeling tool, regardless of calculi.

Problem structuring using influence diagrams implies choices of apparent causality and (conditional) independence between world events. Star (p. 323), Geiger and Pearl (p. 136), and Verma and Pearl (p. 352) examined the significance of these assumptions. Star challenged their validity for real-world problems, and the other authors presented powerful computational methods for exploiting them.

The use of influence diagrams is especially significant because it represents a trend toward tool building in the UAI community. Shenoy and Shafer's results are computational and lead naturally toward a software shell (not yet implemented) for multiple calculi experimentation. Geiger, Verma, and Pearl's works are all oriented toward computations on influence diagrams. Several on-line tools were presented at the workshop, including influence diagram integration with Mac-II hypermedia by Chavez and Cooper (p. 49), which built upon Shachter's David program (available on IBM compatible and Apple Macintosh products) for influence diagram development and execution. Heckerman and Horvitz demonstrated a PC and optical disk-based product called Intellipath. Founded on a Bayesian calculus and full utility theory, it is a teaching and interactive decision-aiding tool for investigating medical pathologies.

Perhaps the emergent theme of this year's workshop was the realization that continuing theoretical development is required to apply probability,

belief, and possibility theories in such AI areas as automated planning, medical diagnosis, and computer vision. This dichotomy arises especially in the combinatorics inherent in UAI techniques: the proliferation of arcs (under reversal) in influence diagram computation, the exponential size of frames of discernment, or the enormous bookkeeping structures required for possible world representations. It is likely that next year's workshop will also stress specific applications of theory to real-world problems and highlight the practical, yet theoretical, developments that are necessarily developed to make real-world applications possible.

Interest in UAI has grown beyond the capacity of the workshop format used in this and previous years. In 1989, we will expand the format to permit more papers (a mix of presentations and poster sessions will be used to accommodate the expected volume). To support our increasing size, we invite corporate sponsorship. We also invite the broader AI community to submit papers discussing sound but novel approaches to unsolved problems in representation and reasoning under uncertainty as well as experiences with applications of current approaches to real-world problems.

This year's program chairman was Ross Shachter, and the general chairman was Tod Levitt. The workshop was sponsored by the American Association for Artificial Intelligence, Advanced Decision Systems, and Knowledge Systems Concepts. For information on next year's workshop, please contact Max Henrion of Carnegie-Mellon University or Ross Shachter of Stanford University. A limited number of proceedings from 1987 and 1988 are available by mail order from Angela Erickson at Advanced Decision Systems, 1500 Plymouth Street, Mountain View, CA 94043-1230. Proceedings from 1985 and 1986 are available in book form from North Holland.

* All references are from the proceedings of this workshop. See Advanced Decision Systems. 1988. Proceedings of the Fourth Workshop on Uncertainty in Artificial Intelligence. Mountain View, Calif.: Advanced Decision Systems.

Proceedings of the

Nonmonotonic Reasoning Workshop

The proceedings of this successful AAAI Workshop held 17-19 October 1984 at the Mohonk Mountain House in New Paltz, New York are still available in limited quantities. Papers in the proceedings emphasize formal approaches to nonmonotonicity.

Topics include:

- Default Reasoning
- Dependency-Directed Inference Systems
- Computational Vision
- Circumscription
- Semantics

The proceedings also include a comprehensive bibliography of literature in the field, compiled by Donald Perlis.

**401 pp., 8 1/2" by 11".
ISBN 0-929280-28-8 \$25.00**