AI Magazine Volume 15 Number 3 (1994) (© AAAI)

The Fourth International Workshop on Nonmonotonic Reasoning

David W. Etherington and Henry A. Kautz

■ The Fourth International Workshop on Nonmonotonic Reasoning brought together active researchers in nonmonotonic reasoning to discuss current research, results, and problems of both theoretical and practical natures. There was lively discussion on a number of issues, including future research directions for the field.

The Fourth International Workshop on Nonmonotonic Reasoning, which was sponsored in part by the American Association for Artificial Intelligence and AT&T Bell Laboratories, was held at Hawk Inn and Mountain Resort in Plymouth, Vermont, from 28–31 May 1992. Attendance was limited to 48 people, including 6 students; 16 attendees were from outside North America. (Extended versions of some of the papers from the workshop were published in a special issue of *Fundamenta Informaticae*, volume 21, numbers 1–2.)

The aim of the workshop was to bring together active researchers in nonmonotonic reasoning to discuss current research, results, and problems of both theoretical and practical natures. Contributions on the following themes were particularly solicited:

Applications of nonmonotonic reasoning: Specific cases were either (1) using a nonmonotonic formalism in an AI or general computer science task, such as qualitative physics, databases, learning, logic programming, diagnosis, or robotics or (2) using a nonmonotonic formalism to analyze and gain insight into (that is, not just model) such a task. Authors were encouraged to stress what the motivation was for the work and what was learned from the experience; descriptions of work that attempted to cross disciplines were encouraged.

Computational tractability: Research in nonmonotonic reasoning deliberately ignored computational concerns for many years. Now, there is a growing interest in tractable realizations of the generally intractable theories of nonmonotonic reasoning. We encouraged papers describing either formal or implemented systems that treat the cost of computation seriously. Possible approaches included limited languages (authors were encouraged to explain why proposed limited languages are of practical interest), approximate reasoning, relevance, resource bounds, and argument-based systems.

Computational issues occupied the lion's share of the workshop, both during the sessions and offline, in marked contrast to earlier workshops.

Formal semantics, closing the loop: Research in nonmonotonic reasoning has given rise to many novel semantic constructions: minimal models for circumscription, perfect models for logic programming, epsilon semantics for probabilistic defaults, and so on. We encouraged papers that start to close the loop between this kind of fundamental work and practice: How do the proposed semantic models help in the creation, verification, or implementation of nonmonotonic theories? What criteria should be used to select one semantic formalism over another?

Submissions to the workshop fell into several categories, including applications of nonmonotonic reasoning to planning and mathematical reasoning; analyses of where the computational difficulties lie within existing theories; techniques for approximate, tractable, nonmonotonic reasoning; applications of algorithms from constraint satisfaction, logic programming, and assumption-based truth maintenance system work to nonmonotonic reasoning; and connections between classical nonmonotonic logics and probability-theoretic logics, conditional logics, and logic programming.

Computational issues occupied the lion's share of the workshop, both during the sessions and offline, in marked contrast to earlier workshops. In part, this situation can be attributed to the growing pressure the field is facing to "put up or shut up"; more importantly, though, it seems to be the result of the area's maturing. Where once it was premature to focus on computational issues because what was to be computed was unclear, there has recently been marked convergence among formalisms. This convergence has yielded enough of a common core that it now makes sense to focus on the problem of how to compute those things that formalisms suggest should follow from nonmonotonic theories.

We do not mean to say that the issue of exactly what nonmonotonic reasoning is has been settled: There are still many dimensions along which nonmonotonic formalisms diverge; it is also not clear that there can ever be a universal nonmonotonic formalism. However, the scope of convergence results linking aspects of various nonmonotonic formalisms (including logic programming) that were presented in recent years (including at this workshop) has been remarkable. Although much basic work remains to be done, the consensus seems to be that there is sufficient common ground to warrant serious computational studies.

A number of papers dealt with (in)tractability. Marco Cadoli and Marco Schaerf (both of Universitá di Roma La Sapienza) and David Etherington and James Crawford (both of AT&T Bell Laboratories) focused on approximation methods for achieving tractable default reasoning at the cost of accuracy. Rachel Ben-Eliyahu and Rina Dechter (both of University of California at Los Angeles [UCLA]) presented reductions of propositional default theories to theories of propositional logic, identifying tractable classes and computational mechanisms for some cases as well as reducing the problem to a form in which constraint-satisfaction problem algorithms can be applied. Both Georg Gottlob (Technische Universitat Wien) and Jonathan Stillman (General Electric Research and Development Center) showed that a variety of propositional nonmonotonic-reasoning tasks are Σ_2^P complete (and, hence, likely harder than NP complete, although likely easier than P-space complete). Unlike many previous complexity results, however, Gottlob and Stillman were able to identify sources of complexity in a way that will be useful in guiding the search for tractable special cases or approximation techniques.

The panel on tractability began with a discussion of the complexity issues in nonmonotonic reasoning, stressing the theme that one of the original promises of default reasoning was efficient inference and asking how we came to have such negative complexity results. Several promising approaches to achieving practical default reasoning were discussed, including both sound and unsound approximations, the use of defaults to guide search, and default reasoning in logic programming using negation as failure as the nonmonotonic operator. Matthew Ginsberg (Stanford University) also argued, however, that worst-case–complexity results are not necessarily relevant to commonsense reasoning because basically all interesting problems are NP hard (at best). His point was that just confirming that this problem is indeed potentially nasty is not really surprising. What makes the problem interesting is the need to solve the average case quickly as well as to somehow cope with the worst case when it comes up.

There was also considerable interest in connections between formalisms and between nonmonotonic formalisms and other work in AI.

There was also considerable interest in connections between formalisms and between nonmonotonic formalisms and other work in AI. In this vein, Fahiem Bacchus (University of Waterloo), Adam Grove (Stanford University), Joseph Halpern (IBM Almaden), and Daphne Kuller (Stanford University) presented a statistical approach that moves toward unifying defaults with statistical reasoning and provides a basis for generalization: Many metareasoning principles, such as preference for more specific facts, seem to fall out directly rather than have to be added explicitly. Moises Goldszmidt and Judea Pearl (both of UCLA) continued their explorations of the connections between probabilistic reasoning and traditional nonmonotonic reasoning, showing how Bayes nets and Markov shielding (a probabilistic analog of stratification) could be used for nonmonotonic reasoning about causality. Craig Boutillier (University of British Columbia) applied conditional logic to unify logics for belief revision and default reasoning, and L. Farinas del Cerro, A. Herzig, and J. Lang (all of Université Paul Sabatier) also connected nonmonotonic and conditional logics. Halina Przymusinska (California State Polytechnic) and Teodor Przymusinski (University of California at Riverside) discussed connections between, and generalizations of, logic programming and default logic.

Application-oriented presentations included work on planning, mathematical reasoning, and extensions to assumption-based truth maintenance systems. Ginsberg and Hugh Holbrook (Stanford University) showed that default reasoning could be used in planning in a way that subsumes the capabilities of hierarchical planning and, in many cases, provides significant computational advantages. This work suggests that there might be many applications in which defaults could be applied without necessitating intractable consistency tests. Leo Bertossi (Pontifica Universidad Catolica de Chile) and Ray Reiter (University of Toronto) showed that standard circumscription provides unintuitive results when applied to problems of generating and reasoning about generic mathematical objects (for example, generalized triangles). They then showed that scoped circumscription can achieve better results, suggesting that the motivation underlying scoped reasoning is more broadly applicable than previously thought.

Olivier Raiman and Johan de Kleer (both of Xerox Palo Alto Research Center) observed that many tasks, including diagnosis and dynamic constraint satisfaction, require forms of nonmonotonic reasoning beyond the capabilities of the assumption-based truth maintenance system. They introduced an extension, the *minimality maintenance system,* that has the same functions and supports the same transactions as the assumption-based truth maintenance system except that the underlying theory is minimized according to a user-defined circumscription policy. They provided algorithms based on computing prime implicates, including an incremental version that supports efficient maintenance of the theory under updates.

Several papers addressed more traditional aspects of nonmonotonic reasoning, including semantics, representational issues, and particular formalisms. Robert Stalnaker (Massachusetts Institute of Technology [MIT]) presented a philosophical study of the whole question of what a nonmonotonic consequence relation is, discussing the conditions that such a relation should meet and the confusions that result from insufficient precision about what notion of inference is being represented. He suggested that such a study is a prerequisite for systematizing and evaluating nonmonotonic theories and motivating unifying frameworks. In a related but somewhat more specialized vein, Michael Morreau (Universitaet Stuttgart) discussed the different results of viewing defaults as expressing norms or simply licensing inferences and showed that the two views are formally incompatible. Phillipe Besnard (IRISA, France) and Torsten Schaub (Technische Huchschule Darmstadt) presented a general possible-world semantics for several variants of default logic. Finally, Fangzhen Lin (University of Toronto) and Yoav Shoham (Stanford University), Kurt Konolige (SRI International), Thorne McCarty (Rutgers University), Camilla Schwind (Faculté des Sciences de Luminy), and Pierre Siegel (Université de Provence), and Mark Boddy (Honeywell Research Labs), Robert Goldman (Tulane University), Keiji Kanazawa (UBC), and Lynn Stein (MIT) described specialized representational problems, techniques, and formalisms.

The panel discussion on sources of defaults and default-related information, such as prioritization or specificity precedence, conditionality or directionality, had mixed results. A variety of sources were cataloged, including statistical data, observed regularities, policies, discourse conventions, representational economies, and normative statements. There was less success, however, in determining how default information should be identified or captured. Although there are obvious ways to apply results from learning research, the questions involved in identifying default information (for example, what level of observed exceptions should lead to rejecting a proposed default?) seem to remain almost wide open.

Finally, several attendees of the First International Workshop on Nonmonotonic Reasoning (New Paltz, New York, 1984) were asked to give their impressions of how the field has progressed in the intervening years. There was general consensus on three main areas for evaluation: progress toward solving problems, impact on AI applications, and impact on other disciplines. Although, admittedly, there have been false starts, it seemed uncontroversial that significant progress has been made on many of the technical problems, especially on mathematical and complexity-theoretic aspects, relationships between various nonmonotonic reasoning schemes, and relationships between nonmonotonic reasoning and probability. Ginsberg suggested that less progress has been made on practical applications of nonmonotonic reasoning, connections outside AI, and the finding of general principles of nonmonotonic reasoning.

Benjamin Grosof (IBM T. J. Watson) pointed out several areas of progress. He noted that there has been a trend toward exploring potential applications, if not toward fielded applications, and that theoretical bases for Prolog-style logic programs with negation as failure, frame-based AI languages, and truth maintenance systems have been cleaned up. In addition, he noted that the importance of a number of formal ideas, including groundedness, causality, and specificity, have been recognized. Vladimir Lifschitz (University of Texas at Austin) also observed that nonmonotonic reasoning has had a major impact on logic programming and that the favor was now being repaid with recent progress on implementing nonmonotonic reasoning mechanisms using logic programming. It was pointed out that all these were recognized as goals at the New Paltz Workshop.

There have been some disappointments, mostly having to do with the lack of really practical nonmonotonic reasoning systems and the attendant lack of impact on other areas, including practice in building inheritance and frame systems. However, many problems have been discovered with the popular implementation mechanisms, such as shortest-path inheritance, and some practical systems have been implemented based on sound nonmonotonic reasoning principles. Unfortunately, however, the application of the more-principled approaches that have been discovered has generally had to await the discovery of ways to deal with computational issues that are only now being addressed.

Some also expressed the feeling that toy problems hold too much sway, and there was some agreement that it is time to start trying to deal with large-scale applications. Lifschitz, however, argued that this effort is a red herring because small problems continue to highlight important distinctions. He claimed that there seems to be little virtue in obscuring issues by going to complicated examples simply for their own sake.

The next workshop in this series was held 29 May to 1 June at Castle Dagstuhl, near Saarbrücken, Germany, and was chaired by Gerd Brewka (Gessellschaft fuer Mathematik und Datenverarbeitung) and Ginsberg.

David W. Etherington received his Ph.D. from the University of British Columbia in 1986. He was a research scientist in the AI Principles Research Department at AT&T Bell Laboratories from 1986 to 1993 and is currently a research associate professor in the Computational Intelligence Research Laboratory at the University of Oregon. He is also the founding chair of the IFIP Working Group on Knowledge Representation. His current research interests include nonmonotonic reasoning and formal, tractable theories of knowledge representation.

Henry A. Kautz is a research scientist at AT&T Bell Laboratories. He received his Ph.D. from the University of Rochester in 1987 and in 1989 was presented with the Computer and Thought Award from the International Joint Conferences on Artificial Intelligence. His research includes work on planning, plan recognition, temporal reasoning, and the complexity of default reasoning. His most recent work is on satisfiability testing and software agents.