Fundamenta Informaticae 21 (1994) 1–5
IOS Press

INTRODUCTION TO THE SPECIAL ISSUE

David ETHERINGTON, Henry KAUTZ
AT&T Bell Laboratories, 600 Mountain Ave., Murray Hill, N.J. 07974

Kurt KONOLIGE
Artificial Intelligence Center, SRI International, Menlo Park, CA 94025

In a continuing tradition, the Fourth International Workshop on Nonmonotonic Reasoning gathered together some 45 of the leading researchers in the field for several days at the end of May, 1992. The isolated, countryside atmosphere at Hawk Mountain Inn in Plymouth, Vermont, promoted interaction and fellowship among the group. Not only was the quality of presented papers high, but there was ample opportunity to meet and talk while hiking the grounds surrounding the Inn.

The field of nonmonotonic reasoning is now over 13 years old, tracing its origins to the famous Volume 13 (1–2) of the Artificial Intelligence Journal, in which the formalisms of Circumscription, Default Logic, and Nonmonotonic Modal Logic were all introduced. Since that time the amount of material devoted to nonmonotonic reasoning has continued to grow. The papers of this volume are a cross-section of the best material presented and debated at the Workshop. They represent current ideas and research directions in Nonmonotonic Reasoning, and reflect the growing interaction with related fields, notably Mathematical Logic and Computer Science.

Broadly speaking, these papers can be divided into four main categories:

- Logical foundations (1 paper)
- Semantics of particular formalisms (4 papers)
- Proof theory and computational complexity (1 paper)
- Applications (3 papers)

It is especially pleasing to see the emphasis put on applications of nonmonotonic formalisms. The main success of nonmonotonic formalisms has been their explication of informal methods with a long provenance in AI (a point made in Raymond Reiter's invited talk at the Knowledge Representation Conference in Boston, October, 1992). The more direct application of these formalisms to solve interesting representational problems is a further measure of their importance.

Logical Foundations

The first paper of this volume, Robert Stalnaker's "What is a nonmonotonic consequence relation?", explores the logical underpinnings of nonmonotonic consequence. Stalnaker approaches the general enterprise of nonmonotonic consequence by looking at it from two conceptually divergent points of view: as a means of explicating implicit information contained in a message, and as a type of inductive inference. The basic claim is that various metatheoretic principles of nonmonotonic consequence, such as rational monotony, can be better understood and evaluated in the context of such general underpinnings.

Stalnaker suggests that one intuition about nonmonotonic reasoning comes from the concept of implicit content. As in conversational implicature, more is meant than what is implicitly said, and the implicit content of a message can change nonmonotonically as the message is expanded. On this reading, nonmonotonic consequence has the form:

$$o(p) \vdash q \,,$$

where p is the explicit message, $o(\cdot)$ is an operator that produces the full content of the explicit message, and q is a deductive consequence of $o(p)$.

Looking at nonmonotonic consequence in this way leads to some immediate judgments about metatheoretic principles: for example, right weakening is admissible because q is a deductive consequence of $o(p)$. Stalnaker then goes on to briefly comment on default reasoning, claiming that if $o(\cdot)$ has the structure of default reasoning, then it should obey cautious monotony and cut, two well-known principles.

The autoepistemic logics of Moore and Halpern and Moses are also examined from the implicit content point of view. For these logics, Stalnaker states that there is a confusion between truth in a single model *vs.* validity in all models. Later, arguing from the implicit content view, he arrives at a model-theoretic version of autoepistemic reasoning, which relates to the possible-worlds semantics of Moore or Levesque.

The second method of comprehending nonmonotonic reasoning is based on inductive inference. Here reasoning is considered to be belief revision based on total evidence: q follows from p if an agent is disposed to accept q on learning p (as total evidence). Stalnaker treats Gärdenfors and Makinson's theories here, and considers Kraus, Lehman and Magidor also to be in the same business of belief revision. On his account, accepting one of the metatheoretic principles amounts to choosing among differing belief revision strategies.

The refreshing aspect of Stalnaker's view is that he tries to understand nonmonotonic logics in service of a particular task such as belief revision, rather than as an isolated mathematical system. Judgments about the suitability of a logic are more readily understood and accepted in this way, and advocates of particular systems can suggest concrete arguments supporting their positions.

Semantics

There are four papers in this section, all of which examine alternative semantics for Default Logic. In the first, "Investigations of model-preference defaults," by Boddy, Goldman, Kanazawa, and Stein, examines the model-preference default (MPD) logic of Kautz and Selman, comparing it to circumscription and default logic, and applying it to several domains. MPD initially was viewed as a restricted form of nonmonotonic inference with good computational properties. However, the authors find that the single complete model semantics of MPD makes it unsuitable for many applications.

For example, in encoding inheritance networks with exceptions into MPD, MPD finds only a single maximal model, which doesn't generate the inferences that are normally wanted.

The second paper is "Possible worlds semantics for default logics," by Besnard and Schaub. This paper provides a uniform Kripke semantics for several types of default logics, especially those that are more constrained in their consistency checks than Reiter's original logic.

The Kripke mechanism is an adaptation of a previous semantics of Schaub's, called focused model structures. In this semantics, two classes of first-order interpretations were used: one for validity of the prerequisites of rules, and the other for satisfiability of the justification and consequent. The default rules then define a preference relation among the model structures, in the same manner as Etherington's semantics. The additional class allows a separate, more strict interpretation of consistency among the justifications.

Besnard and Schaub use this basic idea, but shift it to the more general setting of modal logic and Kripke structures. Their technique is to convert a focused model into a class of Kripke structures, with the first focused model class corresponding to the actual worlds, and the second to the accessible worlds. As in the focused model, a set of defaults defines an ordering among Kripke structures. The differences between Reiter default logic and constrained default logic emerges as a difference in how the defaults generate an ordering.

In the third paper, "Stationary default extensions," Przymusinska and Przymusinski continue their previous work on stationary semantics for autoepistemic logic and logic programming, and apply it to default logic. The basic idea is to use fixed points of the square of Reiter's original Γ operator. The resulting stationary extensions have some interesting properties: they generalize the concept of a default extension; and there is always a least stationary extension, which is the intersection of all stationary extensions. The least stationary extension obeys the property of cumulative monotony.

Finally, "A modal logic for hypothesis theory" (Schwind and Siegel) presents Hypothesis Theory, a system similar to default logic. This theory is a bimodal system, with operators Lp (for believed propositions) and Hp (for hypothesized propositions). The connection between the operators is that $Hp \models \neg L \neg p$, that is, something that is hypothesized cannot be disbelieved. The basic idea of the logic is to form extensions by maximizing a consistent set of hypotheses formulas Hp, relative to the premises. Correspondences between default logic and autoepistemic logic are shown for rational extensions, those for which $\neg Hp \models L\neg p$ holds for every p.

Complexity

In their paper, "Approximate inference in default logic and circumscription," Cadoli and Schaerf extend their approximation technique for classical propositional logic to default logic and circumscription. The bases of the approximation are two inference relations, each parameterized by a set S of sentences. Sentences within S receive a classical two-valued interpretation, while those in its complement get a 3-valued (for the sound but incomplete relation) or a 1-valued (for the complete but unsound relation) interpretation. In this way the classical consequence operation is approached from above and below as the set S grows larger.

To extend this method to nonmonotonic reasoning, the authors note two sources of complexity. One is the choice of generating defaults for an extension, and the second is the consistency check required for their justifications. The authors attack both of these simultaneously, using the two approximate consequence operators. Depending on how they are employed in re-defining the logic, they arrive at either a sound but incomplete system, or a complete but unsound one. These approximations have polynomial complexity for normal propositional default logic (given fixed S).

Similar considerations apply to their treatment of circumscription, which is double exponential in the propositional case. Again, the authors arrive at two approximations, both of polynomial complexity (given fixed S).

Applications

The applications in this volume show the broad reach of nonmonotonic reasoning. They range from formalization of very abstract mathematical practice dealing with generic objects (Bertossi and Reiter), to theories of permission and obligation (McCarty), to the concrete business of building hierarchical planners (Ginsberg and Holbrook).

The first paper, "Circumscription and Generic Mathematical Objects" (Bertossi and Reiter), proposes a technique for describing the informal mathematical practice of reasoning with generic objects. They start with the example of generic triangles, as used in geometric proofs. The concept of genericity is relative to properties of interest: for example, if the properties are same-length sides and right angles, then generic triangles will be neither isosceles or right-angled.

The straightforward way to generate generic triangles is to use circumscription to minimize the appropriate properties. The authors show that five natural minimization encodings do not work, in that they do not generate non-isosceles, non-right-angled triangles. Curiously enough, they are able to show a kind of lottery paradox for these minimizations. Given the connection to the lottery paradox, it is natural to use scoped circumscription as the minimization technique, since it was developed as a way of solving the lottery paradox. The main idea is to restrict the scope of minimization to just the generic individuals.

The beauty of this paper lies in the broad applicability of the result. Although the triangle example could be trivially solved by writing down appropriate restrictions on generic elements, in general it is not easy to consistently define generic objects by inspection, especially when the constraints are complicated. The method should also have application beyond pure mathematical systems, for example, in layout design where one seeks the "most general" position for a new element.

The second paper, "Defeasible deontic reasoning" by L. Thorne McCarty, presents a means of analyzing deontic reasoning through the combination of a deontic logic with a system of defeasible reasoning, and illustrates the method with an example of Chisolm's paradox.

Deontic logic is concerned with questions of obligation and permission: what an agent ought to do, and what he is forbidden to do. Typically, one would like to make local statements about obligation: in situation S, one is obligated to do action A. An agent's obligatory or forbidden actions are then determined by considering this set of statements as a whole. This paradigm parallels that of epistemic reasoning in the face of defaults, in which individual defaults are stated in a modular fashion, and the set of beliefs is determined by maximizing the set of defaults that apply. As in the epistemic case, modular obligations can interfere when considered together, and a straightforward mapping to deontic logic can lead to inconsistencies, which is the source of Chisolm's Paradox.

Since default logic was developed to recover consistency in the case of beliefs, it might be suspected that it can be used in the deontic case, so that statements about obligation are taken to be defeasible. Horty has pursued this line for the case of unconditional obligations, using a normal default logic. However, he does not consider the case of conditional obligations, in which the agent's duty is conditioned on the state of affairs. McCarty tackles this more complex case, using his own brand of deontic logic and defaults with explicit exceptions.

The final paper, "What defaults can do that hierarchies can't" (Ginsberg and Holbrook), is a unique application of default reasoning to the problem of planning. There is a split in planning work between practical planners that use operator descriptions, such as STRIPS and SIPE, versus declarative axiomatizations of action that have their basis in McCarthy and Hayes' Situation Calculus. The latter, while useful as an analytic tool, has not yet been made into a practical planning system. Part of the problem is the lack of a hierarchical representation: typical operational planners solve an easy version of the planning problem at a very abstract level, and then refine the plan to more specific details.

The focus of this paper is representational: showing how the declarative approach can incorporate the advantages of abstraction hierarchies, through the clever use of defaults. The basic idea is simple: provide information about the relative ease of achieving various preconditions to action. This information, expressed axiomatically, furnishes a method for forming hierarchies dynamically. For example, suppose a predicate p can be achieved with effort 1, and q with effort 2, so that p is "easier" to achieve. A plan at a very abstract level would ignore both p and q whenever they were subgoals, assuming they could be achieved. A more precise plan would ignore p but try to satisfy q; and the detailed plan would try to satisfy both.

The authors show that, with this technique, declarative formalisms can produce exactly the same results as operator systems. In addition, there are some advantages to using the hierarchies generated by defaults, stemming from their dynamic nature. The most compelling of these is that there are some problems for which the static hierarchies of operator systems are ill-suited, notably those in which actions can have significant ramifications. In some cases, the static hierarchies must divide actions into an exponential number of subtypes, while dynamic hierarchies can generate the necessary hierarchy "on the fly." Given these promising theoretical results, the authors speculate that it may be possible to build planning systems using declarative formalisms that are competitive with operator systems.