

Nonmonotonic Logics

Nonmonotonic logics are used to formalize plausible reasoning. They allow more general reasoning than standard logics, which deal with universal statements. For example, standard logics can easily represent the argument:

All men are mortal
Socrates is a man

Therefore, Socrates is mortal

$((x)(\text{Man}(x) \Rightarrow \text{Mortal}(x)); \text{Man}(\text{Socrates}) / \text{Mortal}(\text{Socrates}))$

but cannot represent reasoning such as:

Birds typically fly
Tweety is a bird

Therefore, Tweety (presumably) flies.

Such arguments are characteristic of commonsense reasoning, where *default* rules and the absence of complete information are prevalent.

The most salient feature of nonmonotonic reasoning is that the conclusion of a nonmonotonic argument may not be correct. For example, if Tweety is a penguin, it is incorrect to conclude that Tweety flies. Nonmonotonic reasoning often requires jumping to a conclusion and subsequently retracting that conclusion as further information becomes available. Thus, as the set of assumptions grows, the set of conclusions (theorems) may shrink. This reasoning is called *nonmonotonic* in contrast to standard logic which is *monotonic*: as one's set of assumptions grows, one's set of theorems grows as well. (Formally, a system is monotonic if for any two theories A and B, whenever A is a subset of B, the theorems of A are a subset of the theorems of B.)

All systems of nonmonotonic reasoning are fundamentally concerned with the issue of consistency: ensuring that conclusions drawn are consistent with one another and with the assumptions.

Default Logic (Reiter 1980) introduces a *default rule*, an inference rule consisting of an assumption, an appeal to the consistency of some formula, and a conclusion. For example, the rule *Birds typically fly* could be written as

$$\frac{Bird(x):Fly(x)}{Fly(x)}$$

which reads: if x is a bird, and it is consistent that x flies, then conclude that x flies.

Default rules must be applied with care, since conflicting default rules could cause inconsistency if used together. For example, the default *Quakers are usually pacifists* conflicts with the default *Republicans are usually non-pacifists* in the case of Richard Nixon, who was both a Quaker and a Republican. Applying the first default yields the conclusion that Nixon was a pacifist; applying the second default yields the conclusion that Nixon was a non-pacifist; applying both yields inconsistency. One generates *extensions* of a default theory by applying as many default rules as possible. Multiple extensions, or their equivalent, arise in all nonmonotonic logics. The existence of multiple extensions may be seen as a feature or a problem. On the one hand, they allow the expression of reasonable but conflicting arguments within one system, and thus model well commonsense reasoning and discourse. However, conflicting defaults may give rise to unexpected extensions, corresponding to arguments that seem odd or unreasonable. An example of problematic multiple extensions is the Yale Shooting Problem, discussed below.

Autoepistemic Logic (Moore 1985) formalizes nonmonotonicity using sentences of a MODAL LOGIC of belief with belief operator L . Autoepistemic Logic focusses on *stable sets* of sentences (Stalnaker 1980) — sets of sentences which can be viewed as the beliefs of a rational agent — and the *stable expansions* of a premise set. Properties of stable sets include consistency and a version of negative introspection: if a sentence P does not belong to a belief set, then the sentence $\neg L P$ belongs to the belief set. This corresponds to the principle that if an agent does not believe a particular fact, he believes that he does not believe it. To formalize the Tweety example, one represents the rule that birds typically fly as an appeal to an agent's beliefs:

$$L(Bird(x)) \wedge \neg L(\neg Fly(x)) \Rightarrow Fly(x)$$

If I believe that x is a bird and I don't believe that x cannot fly, then (I will conclude that) x flies. Any stable expansion of the premise set consisting of this premise and the premise that Tweety is a bird will contain the conclusion that Tweety flies.

Circumscription (McCarthy 1980;1986) seeks to formalize nonmonotonic reasoning within classical logic by circumscribing, or limiting the extension of, certain predicates. The logic limits the objects in a particular class to those that *must be* in the class. For example, consider the theory containing assumptions that typical birds fly, atypical (usually called *abnormal*) birds do not fly, penguins are atypical, Opus is a penguin, and Tweety is a bird. Opus *must be* in the class of atypical, non-flying birds, but there is no reason for Tweety to be in that class; thus we conclude that Tweety can fly. The circumscription of a theory is achieved by adding a second-order axiom (or, in a first-order theory, an axiom schema), limiting the extension of certain predicates, to a set of axioms.

The systems above describe different ways of determining the nonmonotonic consequences of a set of assumptions. **Entailment Relations** (Kraus *et al.* 1990) generalize

these approaches by considering a nonmonotonic entailment operator $|\sim$, where $P |\sim Q$ means that Q is a nonmonotonic consequence of P , and by formulating general principles characterizing the behavior of $|\sim$. These principles specify how $|\sim$ relates to the standard entailment operator \vdash of classical logic, and how meta-statements referring to the entailment operator can be combined.

Belief Revision (Alchourron *et al.* 1985) studies nonmonotonic reasoning from the dynamic point of view. It focusses on how old beliefs are retracted as new beliefs are added to a knowledge base. There are four interconnected operators of interest: contraction, withdrawal, expansion, and revision. In general, revising a knowledge base follows the Principle of Minimal Change: one conserves as much information as possible.

Nonmonotonic reasoning systems are useful only if they can be successfully **integrated** with other theories of commonsense reasoning. Attempts at integration are often surprisingly difficult. For example, the Yale Shooting Problem (Hanks & McDermott 1987) showed that integrating nonmonotonic reasoning with TEMPORAL REASONING was complicated by the multiple extension problem. The Yale Shooting Problem consists of determining what happens to a turkey when we know that

- (1) a gun is loaded at 1:00 and fired at 3:00;
- (2) firing a loaded gun at a turkey results in the turkey's immediate death;
- (3) guns typically stay loaded (default rule D1);
- (4) turkeys typically stay alive (default rule D2).

Two extensions arise, only one of which is expected. In the expected extension, D1 is applied, the gun remains loaded at 3:00, and the turkey dies. In the unexpected extension, D2 is applied and the turkey is therefore alive after 3:00. This entails that the gun mysteriously became unloaded between 1:00 and 3:00. The problem of formalizing a system of temporal reasoning so that these unexpected extensions do not arise has become a central topic in nonmonotonic research. In nonmonotonic temporal reasoning, it has resulted in the development of theories of CAUSATION and EXPLANATION. (Morgenstern 1996 and Shanahan 1997 give summaries and analyses.)

Integrating nonmonotonic logics with MULTIAGENT SYSTEMS is also difficult. The major problem is modelling nested nonmonotonic reasoning: agents must reason about other agents' nonmonotonic reasoning processes. Most nonmonotonic formalisms are not expressive enough to model such reasoning. Moreover, nested nonmonotonic reasoning requires that agents know what other agents do *not* believe, a difficult requirement to satisfy (Morgenstern & Guerreiro 1993).

In general, integration may require extending both the nonmonotonic formalism and the particular theory of commonsense reasoning.

Applications of nonmonotonic systems are scarce. Implementors run into several difficulties. First, most nonmonotonic logics explicitly refer to the notion of consistency of a set of sentences. Determining consistency is in general undecidable for first-order theories; thus predicate nonmonotonic logic is undecidable. Determining inconsistency is decidable but *intractable* for propositional logic; thus performing propositional nonmonotonic reasoning takes exponential time. This precludes the development of general efficient nonmonotonic reasoning systems (Selman & Levesque 1993).

However, efficient systems have been developed for limited cases. LOGIC PROGRAM-

MING, the technique of programming using a set of logical sentences in clausal form, uses a nonmonotonic technique known as “negation as failure” : a literal, consisting of an atomic formula preceded by a non-classical negation operator, is considered true if the atomic formula cannot be proven. Although logic programs cannot express all types of nonmonotonic reasoning, they can be very efficient (Gottlob 1992). Likewise, inheritance with exceptions is an efficient, although limited, form of nonmonotonic reasoning (Horty *et al.* 1990, Stein 1992). These limited cases handle many common types of nonmonotonic reasoning.

Due to its efficiency, logic programming has been used for many applications, ranging from railway control to medical diagnosis (*Proceedings of the Conference on Practical Applications of Prolog* 1996;1997), although few applications exploit logic programming’s nonmonotonic reasoning abilities. Aside from logic programming, nonmonotonic logic is still rarely used in the commercial world. This may be because the nonmonotonic reasoning community and the commercial world focus on different problems, and because there are few industrial-strength nonmonotonic tools (Morgenstern 1998).

There are similarities between nonmonotonic logics and other areas of research which seek to formalize UNCERTAIN REASONING, such as FUZZY LOGIC and PROBABILISTIC REASONING (especially BAYESIAN NETWORKS). These fields are united in their attempt to represent and reason with incomplete knowledge. The character of nonmonotonic reasoning is different in that it uses a qualitative rather than a quantitative approach to uncertainty. Attempts to investigate the connections between these areas include the work of Goldszmidt and Pearl (1993;1996).

—Leora Morgenstern

See also

COMPUTATIONAL COMPLEXITY
 FRAME PROBLEM
 KNOWLEDGE REPRESENTATION
 PROBABILITY, FOUNDATIONS OF

REFERENCES

- Alchourron, C.E., P. Gardenfors, and D. Makinson. (1985). On the logic of theory change: partial meet functions for contraction and revision. *Journal of Symbolic Logic* 50:510-530.
- Goldszmidt, M. and J. Pearl. (1992). Rank-based systems: a simple approach to belief revision, belief update, and reasoning about evidence and actions. *Principles of Knowledge Representation and Reasoning: Proceedings of the Third Conference (KR-92)*.
- Goldszmidt, M. and J. Pearl. (1996). Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence*.

- Gottlob, G.. (1992). Complexity results for nonmonotonic logics. *Journal of Logic and Computation* 2(3):397-425.
- Hanks, S. and D. McDermott (1987). Nonmonotonic logic and temporal projection. *Artificial Intelligence* 33(3):379-412.
- Horty, J., R. Thomason, and D. Touretzky. (1990). A skeptical theory of inheritance in nonmonotonic semantic networks. *Artificial Intelligence* 42:311-49.
- Kraus, S., D. Lehmann, and M. Magidor. (1990). Nonmonotonic reasoning, preferential models, and cumulative logics. *Artificial Intelligence* 44:167-207.
- McCarthy, J. (1980). Circumscription—a form of nonmonotonic reasoning. *Artificial Intelligence* 13:27-39.
- McCarthy, J. (1986). Applications of circumscription to formalizing common-sense knowledge. *Artificial Intelligence* 28:86-116.
- Moore, R. (1985). Semantical considerations on nonmonotonic logic. *Artificial Intelligence* 25(1):75-94.
- Morgenstern, L. (1996). The problem with solutions to the frame problem. In K. Ford and Z. Pylyshyn (Eds.), *The Robot's Dilemma Revisited*. Ablex, pp. 99-133.
- Morgenstern, L. (1998). Inheritance comes of age: applying nonmonotonic techniques to problems in industry. *Artificial Intelligence* 103. Shorter version in *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI-97)*, Los Altos, Morgan Kaufmann, pp. 1613-1621.
- Morgenstern, L. and R. Guerreiro. (1993). Epistemic logics for multiple agent nonmonotonic reasoning I. *Proceedings of the Second Symposium on Logical Formalizations of Commonsense Reasoning (CS-93)*
- Proceedings of the Fourth Conference on Practical Applications of Prolog.* (1996).
- Proceedings of the Fifth Conference on Practical Applications of Prolog.* (1997).
- Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence* 13:81-132.
- Selman, B. and H. Levesque. (1993). The complexity of path-based defeasible inheritance. *Artificial Intelligence* 62(2):303-340.
- Shanahan, M. (1997). *Solving the Frame Problem*. Cambridge: MIT Press.
- Stalnaker, R. (1980). A note on nonmonotonic modal logic. Department of Philosophy, Cornell University.
- Stein, L. (1992). Resolving ambiguity in nonmonotonic inheritance hierarchies. *Artificial Intelligence* 55:259-310.

Further Readings

- Antoniou, G. (1997). *Nonmonotonic Reasoning*. Cambridge: MIT Press.
- Besnard, P. (1989). *An Introduction to Default Logics*. Berlin:Springer.
- Brewka, G. (1991). *Nonmonotonic Reasoning: Logical Foundations of Commonsense*. Cambridge: Cambridge University Press.
- Dix, J., U. Furbach, and A. Nerode (Eds.). (1997). *Logic Programming and Nonmonotonic Reasoning*, Berlin: Springer.
- Etherington, D. (1988). *Reasoning with Incomplete Information*. London: Pitman.

- Gabbay, D., C.J. Hogger, and J.A. Robinson (Eds.). (1994). *Handbook of Logic in Artificial Intelligence and Logic Programming, vol. 3: Nonmonotonic Reasoning and Uncertain Reasoning*. Oxford: Clarendon Press.
- Gardenfors, P. (1988). *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Cambridge: MIT Press.
- Gardenfors, P. (Ed.). (1992). *Belief Revision*. Cambridge: Cambridge University Press.
- Geffner, H. (1992). *Default Reasoning: Causal and Conditional Theories*. Cambridge: MIT Press.
- Genesereth, M. and N. Nilsson. (1987). *Foundations of Artificial Intelligence*. San Mateo: Morgan Kaufmann.
- Ginsberg, M. (Ed.). (1987). *Readings in Nonmonotonic Reasoning*. San Mateo: Morgan Kaufmann.
- Konolige, K. (1988). On the relation between default and autoepistemic logic. *Artificial Intelligence*. 35:343-382.
- Lloyd, J. (1987). *Foundations of Logic Programming, 2nd Edition*. Springer.
- Marek, W. and M. Truszczyński. (1993) *Nonmonotonic Logic*. Springer.
- McDermott, D. (1982). Non-monotonic logic II: non-monotonic modal theories. *Journal of the Association for Computing Machinery* 29:33-57.
- McDermott, D. and J. Doyle (1980). Non-monotonic logic I. *Artificial Intelligence* 25:41-72.
- Pearl, J. (1990). Probabilistic semantics for nonmonotonic reasoning: a survey. In G. Shafer and J. Pearl (Eds.), *Readings in Uncertain Reasoning*. Morgan Kaufmann.
- Poole, D. (1988). A logical framework for default reasoning. *Artificial Intelligence* 36:27-47.
- Shoham, Y. (1988). *Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence*. Cambridge: MIT Press.
- Williams, M. (1995). Theory base change: a computational model. *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95)*. Los Altos: Morgan Kaufmann.