

WHAT ELSE IS WRONG WITH NON-MONOTONIC LOGICS?
Representational and Informational Shortcomings

Jane Terry Nutter
Department of Computer Science
State University of New York at Buffalo
Amherst, New York 14226

ABSTRACT

Non-monotonic logics have been used recently for a variety of A.I. purposes, including belief revision and default reasoning in question-answering and expert systems. This paper argues that by their nature, such systems discard information which has a role in human belief systems. In particular, systems which use non-monotonic reasoning lose the distinction between fully justified inferences and reasonable presumptions, in the process losing the ability to record failed expectations as such, an ability which provides a useful measure of salience for A.I. systems.

INTRODUCTION

Many A.I. systems deal less with knowledge than with beliefs which are incomplete, which change, and which frequently include generalizations which are known sometimes to fail. Consequently, it sometimes seems not only desirable but necessary to draw conclusions which are not strictly entailed by the information in the system and which further information might counterindicate. In traditional logics, if a set of premises permits a conclusion to be inferred, any larger set containing the original premises permits the same inference. Logics having this property are called monotonic. Because reasoning from beliefs seems not to share this property, researchers have devised non-monotonic logics for use in various A.I. systems.

Non-monotonic logics have generally been founded either on principles from modal logic (see e.g. [McDermott and Doyle 1980] and [McDermott 1982]) or on Zadeh's work [Zadeh 1965; Zadeh 1968] on fuzzy sets and fuzzy logic (see e.g. [Aronson, Jacobs and Minker, 1980]). A substantial technical literature now provides model theories for and proves theorems about such logics (see e.g. [Lee 1972] and [Reiter 1980]), and several A.I. systems have implemented some degree of non-monotonic reasoning, with or without some additional reasoning scheme (see e.g. [Duda, Hart, Nilsson and Sutherland 1978]). Not all responses have been favorable. Non-monotonic logics have been criticized both for technical shortcomings (e.g. [Davis 1980]) and on more general philosophical grounds, which question using logic to govern inferences which involve complex judgements of causal connections and the like (see e.g. Israel 1980)).

The criticisms I want to put forward here fall into neither of these categories. Instead they have to do with the precise role such logics play in A.I. systems when they are used to model reasoning from default generalizations. I will argue that non-monotonic logics do not and in principle cannot accurately reflect such reasoning. Obviously, this criticism depends strongly on what is to be modeled. Hence the first step is to get clear about what uses of non-monotonic logic I am criticizing. The second section discusses this issue. In the third section, I take up the particular aspects of knowledge and information based on default reasoning which I claim non-monotonic logics by their nature cannot reflect.

THE PHENOMENA

It is not always immediately clear what a proposed system is intended to model. Fuzzy logic in particular has been associated with many different purposes: modelling inference patterns, associating probability measures with events or confidence measures with propositions (neither of which is the same as modelling inference patterns), measuring class membership ("To what extent can we call a bacteriophage an animal?"), and almost anything else which someone might want to measure using numbers between zero and one. Each of these uses raises questions of appropriateness; certainly I don't mean to take them all on here.

I am interested specifically in the use of non-monotonic logics to implement inference in systems which are intended either to model some belief space or to perform as expert systems with extended capacities that would include, for instance, formulating descriptions of things they know about (including facts they learn), explaining what makes some object or event unusual, and so on. The sorts of systems I am talking about can receive new information, store representations of it, perform inferences using new information, and report the results of all these operations in some reasonable fashion. They also perform reasoning based on generalizations in situations of incomplete information ("default reasoning"). At least in principle, they can represent any information in their domains in which a human might be interested, and can form propositions expressing that information (though of course not necessarily the same sentences that a human agent would use). By "in principle" I mean that while they may currently lack these abilities, the abilities are desirable, and it is anticipated that they could be added to the system, although it might take some research to find out exactly how.

Such systems could in principle use non-monotonic logic many different ways. This paper deals with using non-monotonic logic as the basic inferential mechanism for such systems. My claim is this: non-monotonic logics cannot support the capabilities listed above. Moreover, this inability does not stem from any technical feature of existing non-monotonic logics. Rather it results directly from non-monotonicity itself when applied to default reasoning.

THE SHORTCOMINGS

1. Their identity. Suppose we start in out in situation S , with information (including unqualified beliefs and generalizations) supporting but not strictly entailing conclusion C . For instance, we know what birds are, and we know that Roger is a living unplucked bird. This information supports the conclusion that Roger flies, but we know this might fail. Now suppose that we learn something which contradicts our previously supported conclusion. In our example, we might learn that Roger is a kiwi. Call the new situation S^* . According to non-monotonic logic, S entails C , S does not entail not- C , S^* entails not- C , and S^* does not entail C . This set of entailment relations provides a consistent basis for stating whether or not, to the best of our current knowledge, Roger flies. But it is not a sufficient basis for all of our relevant knowledge in either S or S^* .

In situation S , a careful speaker would not say "Roger flies," but something like "It is reasonable to suppose that Roger flies." This qualification would not be placed on all conclusions about Roger: "Roger has feathers" is justified absolutely by our beliefs in S . We may be mistaken that Roger is a live unplucked bird, but if we are right about that, we must be right about his having feathers, since by the biological definition of birds, they all without exception have a genetic disposition to produce feathers.

The relationship between our knowledge and the conclusion that Roger has feathers is fundamentally different from the relationship between our knowledge and the conclusion that Roger flies. Furthermore, people are frequently interested in this kind of difference: questions like "Are you sure?" would not otherwise be so common. There is an important distinction between expectations and knowledge. But non-monotonic logic treats the relationship in both cases as entailment. Worse, it produces exactly the same conclusion in situation S as it would in situation S**, where we add the information that Roger is a mature parrot, uncrippled and in normal health. But we know that Roger-the-parrot flies, whereas we only suspect that Roger-the-bird flies. We already know in situation S that this uncertainty exists. This is the first shortcoming of non-monotonic logic: it loses the distinction, present in the "real life" situation, between justified beliefs and justified assumptions. (For more on this distinction, see [Nutter 1982].)

In situation S*, more information is lost. When we learn that Roger is a kiwi, we don't just know that he doesn't fly. We know that (a) Roger is a bird; so (b) there is reason to believe that Roger flies; but in fact (c) Roger is a kiwi, so (d) he doesn't fly after all. Taken together, (b) and (d) contain substantial information: they tell us that an expectation has failed. Non-monotonic logic forces the difference between justified belief and justified assumption into the logic, so that "There is reason to suppose that Roger flies" becomes the same proposition as "Roger flies"; thus viewed, (b) and (d) represent an out-and-out contradiction. To maintain consistency, (b) must be rejected: as stated above, S* entails not-C ("Roger does not fly") and does not entail C ("Roger flies" — in this system, the same as "There is reason to suppose that Roger flies").

So in situation S*, non-monotonic logics can not support inferring "There is reason to suppose C, but not C". This is information in which a human might be interested. Notice that this problem arises explicitly from non-monotonicity. The information lost in S* is precisely what is known in S. Unless that information is lost, adding premises cannot cause conclusions to be rejected, and the logic is by definition monotonic.

Hence we have two kinds of information which non-monotonic logics fail to support. First, they lose the distinction at the propositional level between knowledge and supposition. Associating a measure instead of a truth value with propositions cannot do the work needed here, especially if the metric is taken as measuring probability or confidence: the probability that a single fair toss of an unbiased coin will land heads is 0.5; and of that fact we are completely certain. Knowledge of probabilities is itself knowledge which may be either justified by other knowledge or only suggested. Furthermore, default generalizations are not statistical [Nutter 1983]. Second, because non-monotonic logics fail to support the distinction between knowledge and supposition, they also fail to support reports of failed expectations.

2. Their importance: identifying salience. Consider the following two descriptions of birds.

(a) Felix is a bird who lives in North America. He is under four feet tall, he flies, and he travels slowly on the ground.

(b) Oscar is a bird who lives in Africa. He is over four feet tall, he can't fly, but he travels very rapidly on the ground.

Which of these birds do we know more about? Felix could be almost any North American bird, except a road runner. Oscar is an ostrich, and couldn't be anything else. Yet in both cases, I have simply given the continent they live on, their height (in vague

terms), whether they fly, and their speed on the ground. How do we come to have so much more information about Oscar than about Felix?

When generalizations hold, they don't tell us much. But when they fail, their failure conveys information. Oscar's height, flightlessness, and ground speed are all unusual; together with his African origins, they pin down his species. (If he were Australian, he'd be an emu.)

Saliency has been called the key to a major natural language generation problem: what and how much to say (see e.g. [Conklin and McDonald 1982]). Most techniques for determining saliency depend on either marking particular properties for a class of objects or determining differences between a pair of objects (see the above, [McCoy 1982] and [McKeown 1982]). Neither of these techniques will let a system produce paragraph (b) when asked to describe Oscar but produce only the first sentence of paragraph (a) when asked to describe Felix. However, consider the following rule: if X belongs to kind K, and members of kind K typically have property P but X does not, that is interesting and should be reported. A representation and underlying logic which distinguishes "There is reason to suppose that P" from "P" and allows deducing "Not P, and there is reason to suppose that P" will support this rule for determining saliency. But the logic of such a system will be monotonic.

CONCLUSION

Non-monotonic logics have been motivated largely by the belief that reasoning from default generalizations involves denying old conclusions on the basis of new evidence consistent with all previous premises. I have argued here that this belief arises from failing to distinguish between guarded statements of the form "There is reason to suppose that P" and statements of the form "P". Given this distinction, "There is reason to suppose that P" and "not P" do not contradict one another, so on learning the second we need not reject the first. Hence a monotonic logic providing this distinction can deal with default reasoning. Furthermore, failing to distinguish between these classes of statements causes information to be lost. In particular, two specific gaps appear: (a) the system will not know and be able to report the difference between those conclusions which its premises warrant without reservation and those conclusions which its premises only suggest, and (b) because the system will lose access to the reasonable assumptions when specific information overrides them, it will be unable to detect and state that a reasonable expectation has failed. Consequently, adopting non-monotonic logics may deny A.I. systems access to a simple and useful rule for determining saliency.

ACKNOWLEDGEMENTS

I want particularly to thank Stuart Shapiro and the members of the SNePs research group for their many helpful comments and criticisms.

REFERENCES

1. Aronson, A.R., Jacobs, B.E., and Minker, J. A note on fuzzy deduction. JACM v. 27 (1980) 599-603.
2. Conklin, E.J. and McDonald, D.D. Saliency: the key to the selection problem in natural language generation. Procs. 20th Annual Meeting of the Association for Computational Linguistics (1982) 129-135.

*Ed Conklin
necessary -
non necessary
distinction
be supported*

3. Davis, M. The mathematics of non-monotonic reasoning. *A.I.* v. 13 (1980) 73-80.
4. Duda, R.O., Hart, P.E., Nilsson, N.J., and Sutherland, G.L. Semantic network representations in rule-based inference systems. In: *Pattern-Directed Inference Systems*, D.A. Waterman and F. Hayes-Roth, editors, Academic Press (New York, 1978) 203-223.
5. Israel, D.J. What's wrong with non-monotonic logic? *Procs. First Annual National Conf. on Artificial Intelligence*, American Association for Artificial Intelligence (1980) 99-101.
6. Lee, R.C.T. Fuzzy logic and the resolution principle. *JACM* v. 19 (1972) 109-119.
7. McCoy, K.F. Augmenting a database knowledge representation for natural language generation. *Proc. 20th Annual Meeting of the Association for Computational Linguistics* (1982) 121-128.
8. McDermott, D.V. and Doyle, J. Non-monotonic logic I. *A.I.* v. 13 (1980) 41-72.
9. McDermott, D. Non-monotonic logic II. *JACM* v. 29 (1982) 33-57.
10. McKeown, K.R. The TEXT system for natural language generation: an overview. *Procs. 20th Annual Meeting of the Association for Computational Linguistics* (1982) 113-120.
11. Nutter, J.T. Defaults revisited, or "Tell me if you're guessing". *Procs. Fourth Annual Meeting of the Cognitive Science Society* (1982) 67-69.
12. Nutter, J.T. Default reasoning in A. I. systems. Draft (1983).
13. Reiter, R. A logic for default reasoning. *A.I.* v. 13 (1980) 81-132.
14. Smith, E.E. and Osherson, D.N. Conceptual combination and fuzzy set theory. *Proc. Fourth Annual Conference of the Cognitive Science Society* (1982) 47-49.
15. Zadeh, L.A. Fuzzy sets. *Inf. Control* v. 8 (1965) 338-353.
16. Zadeh, L.A. Fuzzy algorithms. *Inf. Control* v. 12 (1968) 92-102.

