Nonmonotonic Logics: Meaning and Utility

Yoav Shoham

Stanford University
Computer Science Department[1]

Abstract

We propose a unifying framework for nonmonotonic logics, which subsumes previously published systems, and at the same time is very simple. We discuss some of the technicalities of the new general framework, illustrate briefly how some previous systems are special cases of it, and finish an informal discussion of the intuitive meaning of nonmonotonic inferences.

1 Introduction

The last decade or so has seen many logical systems which support so-called *nonmonotonic inferences.* In these formalisms one is allowed not only traditional inferences of classical logic, but also more "speculative" ones. It is often said that in those systems one may "jump to a conclusion" in the absence of evidence to the contrary, or that one may assign formulas a "default value," or that one may make a "defeasible inference." The prototypical example is inferring that a particular individual can fly from the fact that it is a bird, but retracting that inference when an additional fact is added, that the individual is a penguin. This is why such formalisms are called nonmonotonic: a fact entailed by a theory might no longer be entailed by a larger theory. Of course, classical logic is monotonic.

The original and best-known nonmonotonic logics are McCarthy's *circumscription* [10], Reiter's *default logic* [14], McDermott and Doyle's *nonmonotonic logic I* [11], McDermott's *nonmonotonic logic II* [12], and Clark's *predicate completion* [2]. In recent years many more systems have been suggested, and the old ones were further explored. Lifschitz provided new results on circumscription [7]. Further investigations of default logic include Etherington's work [3] and Lukaszewicz' [9]. Moore's *autoepistemic logic* [13] is an adaptation of McDermott's NML II, and a version of it was further investigated by Halpern and Moses [4].

These various formalisms are very different superficially. For example, circumscription amounts to adding a second-order axiom to a given first-order theory. A default theory, on the other hand, contains a collection of *default rules,,* a notion quite outside classical logic, and its meaning is defined using a fixed-point construction which relies on those default rules. McDermott's and Moore's logics are still different, and formulas in those logics contain modal operators, which are meant to capture the notions of consistency and belief, respectively. The nonstandard nature of the various systems and their diversity has made it hard to gain a good understanding of them and to compare among them.

However, the main problem with existing nonmonotonic formalisms is not their overwhelming complexity, as much as it is their limited expressiveness. In particular, they were all shown to fail to capture nonmonotonic *temporal* inferences. The problems were first reported by Hanks and McDermott [5], in response to which several solutions were offered. One such solution was proposed by Lifschitz, and it was to generalize circumscription to so-called p*ointwise circumscription* [9]. My approach has been to construct the logic of *chronological ignorance* [16], and in doing I defined a very general framework for nonmonotonic logics. The purpose of this paper is to present this general framework.

This paper has two distinct parts. The first part is more technical, and consists of the next two sections. The next section introduce the general framework of nonmonotonic logics, the result being a simple system to which on the one hand all existing nonmonotonic logics can be reduced (and thus easily understood and easily compared to one another), and which on the other hand suggests radically new nonmonotonic logics. The section following that briefly shows how some previous nonmonotonic formalisms can be viewed as special cases of the general framework. The second part of the paper is more held at the intuitive level, and consists of the last section. It discusses the intuitive meaning behind nonmonotonic inferences, and argues that the proposed distinction between default inferences and autoepistemic ones should be abolished.

2 Formal construction of nonmonotonic logics

The basic idea behind the construction is the following. In traditional logic, the meaning of a formula is the set of interpretations that satisfy it, or its set of models (where "interpretation" means truth assignment for PC, a first-order interpretation for FOPC, and a (Kripke interpretation, world)-pair for modal logic). One gets a nonmonotonic logic by changing the rules of the game, and focusing on only a subset of those models, those that are "preferable" in a certain respect (these preferred models are sometimes called "minimal models," a term introduced by McCarthy in connection with circumscription). The reason this transition makes the logic nonmonotonic is as follows. In classical logic $A \models C$ if C is true in all models of A. Since all models of $A \wedge B$ are also models of A, it follows that $A \wedge B \models C$, and hence that the logic is monotonic. In the new scheme we

[1]The work described here was curried out when the author was a graduate •tudent at Yale University.

have that $A \models C$ if C is true in all *preferred* models of A, but $A \wedge B$ may have preferred models that are not preferred models of A. In fact, the class of preferred models of $A \wedge B$ and the class of preferred models of A may be completely disjoint.

Many different preference criteria are possible, all resulting in different nonmonotonic logics. Circumscription, default logic and autoepistemic logic all assume implicitly very specific (and different) preference criteria. In [15] I look into other preference criteria that are suitable for temporal reasoning. In this paper I look more closely at the general framework, and at the meaning in it of notions such as satisfiability and entailment.

The following discussion here will be most general, and when I talk of a "standard logic" I allow in principle any logic with the usual compositional, model-theoretic semantics. Rather than give a precise definition of this condition, however, let me in the following assume that by a "standard logic" one means the propositional calculus or first-order predicate logic, either classical or modal (i.e., one of four logics); extension to other cases (such as higher-order logic) will be obvious.[2]

In order to have the following discussion apply uniformly to the classical and modal logics, let me misuse the terminology slightly by calling anything which goes to the left of the $\models$ an *interpretation*. In the classical cases this is what is indeed usually called an interpretation. In the modal cases it is a pair (Kripke structure,world).

The transition to nonmonotonic logic puts into question notions that are well understood in logic, such as satisfaction, entailment, satisfiability, validity, and proof theory. To see what these might mean in the context of nonmonotonic logic, it will be helpful to recall some definitions from the standard case.

Definition 1 (Reminder)
Let $\mathcal{L}$ be a standard logic, and A and B two sentences in $\mathcal{L}$. The fact that an interpretation M satisfies A is denoted by $M \models A$. In this case we say that M is a model of A. A is satisfiable if A has a model. A is valid if A is satisfied by all interpretations. Clearly A is satisfiable iff $\neg A$ is not valid. A entails B (written $A \models B$) if B is satisfied by all models of A, or equivalently, if all models of A are also models of B.

From these definitions, the deduction theorem follows very easily:

Theorem 1 (Reminder)
Let $\mathcal{L}$ be a standard logic, and A, B and C sentences in $\mathcal{L}$. Then $A \wedge B \models C$ iff $A \models B \supset C$.
Clearly all standard logics are monotonic, in the sense that for any A, B and C in the logic, if $A \models C$ then also $A \wedge B \models C$.

Nonmonotonic logics are the result of associating with a standard logic a preference relation on models. More specifically, we make the following definitions. Let $\mathcal{L}$ be a standard logic and A a sentence in it, and let $\sqsubset$ be a strict partial order on interpretations for $\mathcal{L}$. Intuitively, $M_1 \sqsubset M_2$ will mean that the interpretation M_2 is *preferred over* the interpretation M_1. $\mathcal{L}$ and $\sqsubset$ define a new logic $\mathcal{L}_\sqsubset$. I will call such logics *preference logics*. The syntax of $\mathcal{L}_\sqsubset$ is identical to that of $\mathcal{L}$.In order to define the semantics of $\mathcal{L}_\sqsubset$ we must define the notions of satisfaction (of a formula by an interpretation), satisfiability, validity, and entailment.

Definition 2 *An interpretation M preferentially satisfies A (written $M \models_\sqsubset A$) if $M \models A$, and if there is no other interpretation M' such that $M \sqsubset M'$ and $M' \models A$. In this case we say that M is a preferred model of A.*

Clearly, if $M \models_\sqsubset A$ then also $M \models A$.

In the full paper I define the notions of *preferential satisfiability* and *preferential entailment*. Since those are on the one hand somewhat problematic and on the other hand somewhat unimportant, I will proceed directly to the notion of *preferential entailment*, and that has a very intuitive definition:

Definition 3 *A preferentially entails B (written $A \models_\sqsubset B$) if for any M, if $M \models_\sqsubset A$ then $M \models B$, or equivalently, if the models of B (preferred and otherwise) are a superset of the preferred models of A.*

Observation. If $\mathcal{L}$ is a preferential logic, it may contain sentences A and B such that both $A \models_\sqsubset B$ and $A \models_\sqsubset \neg B$. Furthermore, A need not be inconsistent for this to be the case – it is sufficient that it have no preferred models. Indeed, in preferential logic the role played by preferential satisfiability is exactly analogous to that played by satisfiability in standard logics, witness the following obvious lemma:
Lemma 2 *Let $\mathcal{L}_\sqsubset$ be a preferential logic, and A a sentence in it. Then A is preferentially satisfiable iff there does not exists a sentence B in $\mathcal{L}_\sqsubset$ such that both $A \models_\sqsubset B$ and $A \models_\sqsubset \neg B$.*

Definition 4 *$\mathcal{L}_\sqsubset$ is monotonic if for all $A,B,C \in \mathcal{L}$, if $A \models_\sqsubset C$ then also $A \wedge B \models_\sqsubset C$.*

Observation. The above definition is equivalent to saying that a logic is monotonic if all preferred models of $A \wedge B$ are also preferred models of A. Viewing $\mathcal{L}$ as the special case $\mathcal{L}_\phi$, where ϕ is the empty relation, we note that $\mathcal{L}_\phi$ is monotonic.

It is interesting to see what happens to the deduction theorem in light of the new definition of entailment. It turns out that while the theorem is false in general, a weaker version of it still holds, in which the 'iff' is changed to an 'if':

Theorem 3 *Let $\mathcal{L}_\sqsubset$ be a preferential logic, and A, B and C three sentences in it. Then, if $A \wedge B \models_\sqsubset C$, then also $A \models_\sqsubset B \supset C$.*

The converse to this theorem does not hold: if $A \models_\sqsubset B \supset C$ then it does not necessarily follow that $A \wedge B \models_\sqsubset C$. It is not hard to construct a counterexample, for example one in which C is identically false, $A \wedge B$ is *not* identically false, and B is false in all preferred models of A. In fact, it is easy to show that if the converse to the theorem holds then the logic is necessarily monotonic:

[2]When I say 'a modal logic' I deliberately do not commit myself to a particular modal system; that too is not essential here. Again, rather than define the class of all modal systems to which the following applies, let me in the following assume an S5 system whenever a modal logic is mentioned, and leave it to the reader to extend it to other systems.

Theorem 4 Let L_c be a preferential logic. Then the following two statements are equivalent:

1. *For all* $A, B, C \in \mathcal{L}_{\sqsubset}$, *if* $A \models_{\sqsubset} B \supset C$ *then also* $A \wedge B \models_{\sqsubset} C$.
2. $\mathcal{L}_{\sqsubset}$ is monotonia.

3 Special cases

In [15] T discuss previous nonmonotonic systems in some detail, and show how they can be viewed as special cases of the proposed general framework. In particular, I discuss both McCarthy's and Lifschitz' versions of circumscription, Bossu and Siegel's formalism [1], Reiter's default logic, and a version of Moore's autoepistemic logic due to Halpern and Moses [4]. Here I will be able to provide a comparison with only three of those, and very sketchily at that.

3.1 Circumscription

Circumscribing a formula amounts to adding a second order axiom to a theory. McCarthy's original circumscription axiom (and 1 am using Lifschitz' recent reconstruction of it) was:

$$A(P) \wedge \forall p\ \neg(A(p) \wedge p < P)$$

where p is a predicate variable with free variables x, and p<P stands for

$$\forall x(px \supset Px) \wedge \neg\forall x(Px \supset px)$$

This axiom is one way of defining a preference criterion on models, according to which $M_1 \sqsubset M_2$ if

1. For all x, M_1 and M_2 agree on the interpretation off function symbols and all relation symbols other than P,
2. for all x, if $M_2 \models P(x)$ then also $M_1 \models P(x)$, and
3. there exists a y such that $M_1 \models P(y)$ but $M_2 \not\models P(y)$

Other circumscription axioms embody similar preference criteria. To avoid giving the wrong impression, it must be said that the notion of preferred models was implicit in McCarthy's work from the start. In fact, in his original paper he gave a minimality criterion similar to the one stated above, although in subsequent publications the model-theoretic discussion seemed to play a diminishing role. Other researchers too have addressed the model theory of nonmonotonic logics, such as Lifschitz in [8] and Etherington in [3].

My formulation can be viewed as a suggestion to generalize McCarthy's approach in three ways:

1. Start with any standard logic, not necessarily FOPC. For example, I base my formulations on a standard modal logic.
2. Allow any partial order on interpretations, not only the one implied by a particular circumscription axiom. For example, in [15] I suggest a preference criterion that relies on temporal precedence.
3. Shift the emphasis to the semantics, stressing the partial order on models and not the particular way of defining that partial order. In fact, allow the definition of this partial order in any way that leaves no room for ambiguity. The various circumscription axioms, either McCarthy's original ones or Lifschitz's more recent ones, are one way of doing so, and they are most elegant. In my own formulations I have chosen other means of defining preference criteria.

3.2 Bossu and Siegel

At the time this work was conducted I was unaware of related previous work by Bossu and Siegel [1], or 1 would have made an effort to use their terminology where possible. As things are, I renamed some of the concepts they had come up with, and by now I am too fond of my definitions to let go of them. Let me, however, make clear the connection between the two treatments. The summary of it is that they share the basic semantical approach, although there are some minor technical differences between the two, but that Bossu and Siegel thoroughly investigated what turns out to be a very special case of my general formulation. In a little more detail the connection is as follows.

The main part common to both Bossu and Siegel's treatment and my own is the model-theoretic approach, which posits a partial order on interpretations. There are some minor differences in the precise definitions. For example, whereas I defined a to be preferentially satisfiable (or, in their terminology, minimally modelahle) if φ has a preferred model, Bossu and Siegel require in addition that any nonpreferred model of a have a. better model than it which is preferred. Or, as another example, they explicitly reject the definition I chose for preferential entailment (which they call subimplication and denote by I=), since if p is not preferentially satisfiable (i.e., it has no maximally preferred models) then it entails both ϑ and $\neg\vartheta$. I don't view that as a disadvantage, since preferential satisfiability plays a role that is completely analogous to satisfiability. Thus by the same argument one should object to the regular notion of entailment, since inconsistent (i.e., unsatisfiable) theories entail both V and

These are fairly minor differences, and they are overwhelmed by the similarity in the semantical approach to nonmonotonic logics. There is, however, a big difference between the two treatments, and that is in their generality. Whereas 1 allow starting with arbitrary standard logics as a basis, Bossue and Siegel require starting with FOPC. More crucially, whereas I allow any partial order on interpretations, Bossu and Siegel assume one fixed such partial order. As they themselves say,

> The difference between [John McCarthy's] definition and ours is that McCarthy 'minimizes' on some lit erals only, whereas we 'minimize' on every literal.

Circumscription was discussed in the previous subsection. As we have seen, in the simplest version of the logic, preferred models are those in which $P_1(x_1), \ldots, P_n(x_n)$ are true for as few xi's as possible, where the P_i's are predicate symbols specified separately, and x_i are the arguments to P_i. What Bossu and Siegel do is fix the P_i's to be all the predicate symbols that appear in the theory that is being modeled.

3.3 Minimal knowledge

In [4], Halpern and Moses construct a nonmonotonic modal logic, which is a direct adaptation of Moore's autoepistmic logic (which in turn was the outgrowth of McDermott's NML II). In [15] this logic is discussed in more detail; Here let me just provide the preference criterion on (Kripke interpretation, world)-pairs inherent in their logic. According to it, $M_1 \sqsubset M_2$ if (in the following, a base wff is one containing no modal operators):

1. for all base wffs φ, if $M_2 \models \Box\varphi$ then also $M_1 \models \Box\varphi$, and
2. there exists a base wff φ such that $M_1 \models \Box\varphi$ but $M_2 \not\models \Box\varphi$.

Since $\Box\varphi$ is interpreted in the logic as "φ is known," the result is to prefer models in which as few base wffs as possible are known.

4 Two kinds of nonmonotonic inferences?

My treatment of nonmonotonic logics so far has been purely technical. This section is devoted to discussing the intuitive meaning behind nonmonotonic inferences. I will try to clarify some of the confusion arising from the apparent connection between probabilistic information and nonmonotonic inferences, and in the process argue against the suggestion that there are two fundamentally different sorts of nonmonotonic inferences.

The distinction I am alluding to was suggested by Robert Moore in [13]. He contrasts *default inferences*, which are based on some statistical facts (e.g., that most birds can fly), with *autoepistemic inferences*, which are based on the lack of some particular knowledge. Says Moore,

> Consider my reason for believing that I do not have an older brother. It is surely not that one of my parents once casually remarked, "you know, you don't have any older brothers," nor have I pieced it together by carefully sifting other evidence. I simply believe that if I had an older brother I would surely know about it ...This is quite different from a default inference based on the belief, say, that most MIT graduates are oldest sons ...

On the face of it this distinction is quite appealing (certainly 1 was convinced for a while), but upon closer examination it seems to break down completely.

To begin with, one may note that Moore applies his own logic, labelled an autoepistemic one, to the flying birds example, which he himself characterizes as a default case. Furthermore, consider Moore's own older brother example. If one accepts the statement "if I had an older brother then I'd know it," surely one must also accept the statement "if I *didn't* have an older brother then I'd know it." Yet if we adopt this latter sentence rather than the first one, the opposite inference will follow, namely that I have an older brother. On what basis does one prefer the first sentence to the second one, if at all? Notice that if you adopt *both* .sentences, then you end up with two distinct preferred models one in which you have an older brother and know it, and another in which you don't have an older brother and know it - which isn't much help.

Let me suggest a different distinction than the one made by Moore. Rather than distinguish between different kinds of default inferences, one should distinguish between the *meaning* of sentences on the one hand, and the (extra logical) *reason* for adopting that meaning on the other. The meaning, I argue, can be viewed epitemically. The reason for adopting that meaning is computational economy, which often relies on statistical information.

Consider the flying birds example. The *meaning* of "birds fly by default" is that if 1 don't know that a particular bird can not fly, then it can. The computational *reason* for adopting this meaning is that now whenever a bird can indeed fly, we need not mention the fact explicitly either in external communication with other reasoners, or in "internal communication," i.e., thought it will follow automatically. Of course, if we happen to be talking about a penguin, we had better add the knowledge' that penguins cannot fly, or else we will make wrong inferences. In the long run, however, we win: the overwhelming percentage of birds about which we are likely to speak can indeed fly, and so on average this default rule saves us work. If this gain seems small, consider a realistic situation in which we apply thousands and thousands of such rules.

Can we identify a similar rationale behind the rule "by default, I do not have an older brother"? It is less obvious here, which is why the two cases seem superficially different. Yet such motivation must exists, or else one wouldn't prefer this rule to the opposite one. Perhaps the rationale is again a simple counting argument on average a couple has two children, so the speaker has a 50% chance of being the younger one, in which case there is a 50% chance that the older sibling is male. Thus in 75%) of the cases the speaker does not have an older brother, which is not quite as overwhelming as the percentage of flying birds, but still is higher than 50%.[3] Perhaps in fact the motivation for adopting the default rule is more sophisticated, but *some* motivation must exist.

1 am not at all arguing that one makes p true by default just in case p is true most of the time. As 1 have said, the flipside of making a default assumption is the danger of making faulty inferences. For example, if a bird is being discussed and its type is unknown, we will infer that it can fly even though it might turn out to be a penguin. If this seems harmless, think of making the default inference "people you'll meet on the street will not stab you in the hack" in a city in which only 5% of the population are back slabbers. In this case the relatively small chance of being badly hurt seems to outweigh the computational resources needed to reason about individual people on the street, and the discomfort of wearing a steel-plated vest. Notice that if the 5% dropped to 0.00000000005%,, we'd take off the armor and stop looking darkly at passers by. Indeed, that is exactly how we treat the possibility of a nuclear war. Clearly, one must maximize his expected utility when selecting a nonmonotonic theory.

[3] This argument was suggested by Drew McDermott, half in jest, after he was convinrrd of its conclusion

I offer no general guidelines for making such a selection. All I am suggesting here is to separate the two issues, that of defining the meanings of nonmonotonic logics, and that of selecting one.

Notice. The remainder of this section, in which two more arguments are offered in support of the proposed meaning/utility distinction, assumes acquaintance with McCarthy's circumscription, and with modal logics of minimal knowledge. Those were referred to briefly in the previous section, but the reader who is unfamiliar with them may find the following a bit cryptic. In [15] a more detailed discussion of those logics is offered.

The reader may still be bothered by the fact that circumscription involves only classical sentences, and it is not clear how epistemic notions enter into it. It is not hard, however, to convert circumscription into a logic of minimal knowledge. The basic idea is instead of circumscribing a formula $\varphi(x)$, to add the ,'.'.. (For example, in the flying birds case axiom $\forall x\ \varphi(x) \supset \Box\varphi(x)$.
we add the axiom $\forall x\ \mathtt{CANFLY}(x) \supset \Box\mathtt{CANFLY}(x)$.) Since we prefer models in which as few propositional formulas as possible are known, the effect is to have *p* true of as few x's as possible. The natural reading of the axiom $\varphi \supset \Box\varphi$ is indeed "if φ were true then I'd know it." However, if we take the contrapositive form of the axiom, we get the familiar "default rule": $\Diamond\neg\varphi \supset \neg\varphi$, or "if it is possible that *p* is false then it is."

As a final clincher in the argument for the meaning/utility distinction, let me show how this distinction resolves the Lottery Paradox, discussed, for example, in [12]. The paradox is as follows:

> A lottery is held with 1 million participants, including our friend John. The odds of John's winning are so low that we infer by default that he won't. Yet by the same token we can infer that none of the other 999,999 members will win, which contradicts our knowledge that at least one person must win.

In the logic of minimal knowledge (which, as we have just seen, can be translated back into circumscription) we describe the situation bv

$$[\mathtt{WIN}(\mathtt{JOHN}) \lor \mathtt{WIN}(P_1) \lor \mathtt{WIN}(P_2) \lor \ldots \lor \mathtt{WIN}(P_{999,999})] \land$$
$$\forall x\ \mathtt{WIN}(x) \supset \Box\mathtt{WIN}(x)$$

The most ignorant models are ones in which one of the million people wins and we know that he won it, but the identity of that person varies among models. We are therefore not justified in concluding that any particular individual will not win, since there are models in which he does. True, those models are vastly outnumbered by those in which he does not win, but nonmonotonic logics do not let us express the property of a proposition being true in "most" models. This once again shows that such probabilistic information plays a role in choosing the meaning but not in defining it.

So in the above formulation we cannot conclude that John will not win, nor should we want to. If one claims that such a conclusion *is* one that corresponds to default reasoning people use, one must agree to the conclusion that no rational person would ever buy lottery tickets, a prediction that obviously isn't born out in reality.[4] However, the above formulation is not as useless as it might appear. We can still make inferences involving the possibility of people winning, in which rather than condition an inference on their not winning, we condition it on the *possibility* of their not winning. For example, we may add the sentence

$$\Diamond\neg\mathtt{WIN}(\mathtt{ME}) \supset \neg\mathtt{SHOULD\text{-}RESIGN\text{-}JOB}(\mathtt{ME})$$

The rationale here is again statistical. Although there is a model in which I win the lottery and therefore needn't bother teaching a course on AI, it would be foolish for me to resign on that basis.

[4]The argument thai indeed no rational person should, given thai in all lotteries the expected winning amount is negative, is irrelevant. The example would still hold if some bored millionaire organized the lottery, charging each participant one dollar and giving the winner one million and one dollars.

5 Summary

I have presented a uniform approach to constructing and understanding nonmonotonic logics. The value in the formulation has not been its mathematical sophistication, but rather the opposite: the only notion added to traditional logic was that of a partial order on interpretations. The simplicity of the formulations makes transparent what in other systems is less immediate. The formulation is not only simple but also very general, and subsumes previous systems. I briefly indicated how previous systems are a special case of this general framework; in [15], a more detailed comparison with previous nonmonotonic systems is provided. Also, I have argued against the proposed distinction between different kinds of nonmonotonic inferences.

Several open questions remain. One of them has to do with the relation to other nonstandard logical formulations. In particular, Johan van Benthem has drawn my attention to the close parallels between my formulations and formulations in Conditional Logic, as pioneered by D. Lewis [6,18]. In CL one has, instead of a partial order on interpretations, a similarity measure on possible worlds, and the notion of counterfactual entailment is similar to my notion of preferential entailment. I would like to understand this connection better.

Another open question is which particular instances of the general framework I have outlined are interesting from the practical point of view As I have said, the general treatment of nonmonotonic logics offered here grew out of limitations of existing systems. It is still the case, however, that very few concrete preference criteria on models have been investigated. The most common ones are those embodied in the various circumscription axioms and in some default theories. One preference criterion transcending those is that of *chronological ignorance,*, presented in [16]. Lifschitz too has investigated an instance of pointwise circumscription that does not collapse into "old" circumscription. All these, however, are still a drop in the bucket, and we have yet to understand which of the many possible partial orders on models are of practical importance.

Bibliography

1. G. Bossu and P. Siegel, Saturation, Nonmonotonic Reasoning, and the Closed World Assumption, Artificial Intelligence, 25(1), January 1985, pp. 13-65.

2. K. L. Clark, Negation as Failure, in H, Gallaire and J. Minker (Eds.), Logic and Databases, Plenum Press, New York, 1978, pp. 293-322.

3. D. W. Etherington, Reasoning with Incomplete Information: Investigations of Non-Monotonic Reasoning, Ph.D. Thesis, Computer Science Department, University of British Columbia, 1986.

4. J. Halpern and Y. Moses, Towards a Theory of Knowledge and Ignorance: Preliminary Report, Technical Report RJ 4448 48136, IBM Research Laboratory, San Jose, October 1984.

5. S. Hanks and D. McDermott, Temporal Reasoning and Default Logics, Technical Report YALU/CSD/RR 430, Yale University, October 1985.

6. D. Lewis, Counterfactuals and Comparative Possibility, in Harper et al. (Eds.), Ifs, Reidel, Dortrecht, 1973 (book: 1981), pp. 57-85.

7. V. Lifschitz, Computing Circumscription, Proc. 9th IJCAI, August 1985.

8. V. Lifschitz, Pointwise Circumscription, Proc. AAAI, Philadelphia, August 1986.

9. V. Lukaszewicz, Considerations on Default Logic, Proc. AAA I Workshop on Nonmonotonic Reasoning, 1984, pp. 165-193.

10. J. M. McCarthy, Circumscription - a Form of Nonmonotonic Reasoning, Readings in Artificial Intelligence, Tioga Publishing Co., Palo Alto, CA, 1981, pp. 466-472.

11. D. McDermott and J. Doyle, Nonmonotonic Logic I, Artificial Intelligence 13, 1980, pp. 41-72.

12. D. McDermott, Nonmonotonic Logic II: Nonmonotonic Modal Theories, J ACM, 29(1), 1982, pp. 33-57.

13. R. C. Moore, Semantical Considerations on Nonmonotonic Logic, Proc. 8th IJCA1, Germany, 1983.

14. R. Reiter, A Logic for Default Reasoning, Artificial Intelligence, 13, 1980, pp. 81-132.

15. Y. Shoham, Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence, Ph.D. Thesis, Yale University, Computer Science Department, 1986.

16. Y. Shoham, Chronological Ignorance: Time, Nonmonotonicity and Necessity, Proc. AAAI, Philadelphia, August 1986.

17. J. van Benthem, Foundations of Conditional Logic, Journal of Philosophical Logic, 13, January 1984, pp. 303-349.