RC 20124 (07/07/95)
Computer Science

IBM Research Report

Prioritizing Multiple, Contradictory Sources in Common-Sense Learning by Being Told; or, Advice-Taker Meets Bureaucracy

Benjamin N. Grosof

IBM Research Division
T.J. Watson Research Center
P.O. Box 704, Yorktown Heights, NY 10598
(914) 784-7783 direct -7455 fax -7100 main
Internet: grosof@watson.ibm.com (alt.: grosof@cs.stanford.edu)
World Wide Web: http://www.watson.ibm.com

IBM Research Division
Almaden · T.J. Watson · Tokyo · Zurich

This paper appeared in Proceedings of Common Sense '93: *The Second Symposium on Logical Formalizations of Common-Sense Reasoning, edited by Leora Morgenstern* et al. *Held January 11-13, 1993, in Austin, Texas. Proceedings available through the editor leora@watson.ibm.com . This formatted version is available as an IBM Research Report on the World Wide Web http://www.watson.ibm.com .*

Abstract

An important common-sense capability is to learn by "taking advice", or "being told", from a variety of sources of information: messages from various other agents, reading various texts, etc.. Such learned or "assimilated" knowledge provides an important basis for one's actions. Yet a basic feature of life is that one cannot believe everything that one is told. Not only may advice be incorrect; e.g., it may be contradicted by direct experience. Worse, different sources of advice information may contradict each other, or even themselves. How is an advice-taking agent to maintain a consistent, yet usefully actionable, set of beliefs, then? Here, we address this problem, which has not received much previous attention in the AI literature. We illustrate, with an actual example, how it arises in the domain of *bureaucracy*. A common human strategy is to treat advice as working belief, which can be overridden by other advice and direct experience. This suggests the first element of our approach: to represent advice as default-status knowledge.

A crucial issue in logically formalizing common-sense reasoning is then how to represent *precedence* among multiple, potentially contradictory sources of information. We observe that bases for such precedence include reliability and authority. We observe furthermore that precedence relationships may themselves be inferrable and even defeasible. We show how, for the first time, to represent advice-taking, with such precedence among sources, as reasoning in a non-monotonic logical formalism. We employ a generalized variant of circumscription that we have recently developed: Defeasible Axiomatized Policy (DAP) circumscription. (See attached paper detailing DAP circumscription.) Precedence among sources is represented explicitly via the mechanism of prioritization in DAP circumscription. DAP circumscription is the first non-monotonic logical formalism to express *defeasible* prioritization-type precedence.

Our approach to representation of precedence among sources has a number of advantages over other non-monotonic logical formalisms. It treats precedence more explicitly than all previous approaches except perhaps Lifschitz' [1988], who was the first to address the representation of precedence among multiple sources of default information. Compared to Lifschitz', our approach has several advantages. A fundamental advantage is that ours is capable of expressing defeasible prioritization; Lifschitz' is not. Another advantage is that our approach is more adequate and more natural representationally: in particular when, as is common, the precedence partial order is *not* layered ("stratified"). Our approach also represents the concepts of defaults and priorities directly, unlike Lifschitz' which uses only the concept of varying versus not varying. In addition, our approach is better-behaved, e.g., with respect to satisfiability.

Contents

1 Introduction

Note to Reader: Due to constraints of space, as well as considerations of focus, many details about our new "DAP circumscription" formalism (e.g., its precise definition) do not appear in this paper, but are to be found instead in [Grosof, 1993]. That paper thus acts as a companion to this paper.

An important common-sense capability is to learn by "taking advice", or "being told", as proposed, for example, by [McCarthy, 1968] and discussed, for example, in [Dietterich, 1982]. By taking advice, we mean adopting things one is told as one's own beliefs, more or less directly. Advice may arise from a variety of sources of information: messages from various other agents (e.g., an inputting "user"), reading various texts, etc.. Such learned or "assimilated"[1] knowledge provides an important basis for one's actions. Yet a basic feature of life is that one cannot believe everything that one is told. Not only may advice be incorrect; e.g., it may be contradicted by direct experience. Worse, different sources of advice information may contradict each other, or even themselves (e.g., a source may contradict what it earlier told).

In this paper, we address **the problem** of how to handle the potentially contradictory multiplicity of advice. This is an issue that has not received much attention previously in the AI literature.

Humans often employ the **tactic** of treating advice sources as credible, but allowing advice information to be defeasible. In this paper, we show how to formalize this tactic: in our approach, we treat pieces of advice as defaults; more precisely, as default premise beliefs in a variant of the circumscription formalism [McCarthy, 1980] [McCarthy, 1986] [Lifschitz, 1984]. The advice-taking agent then maintains a consistent[2], yet usefully actionable, set of beliefs. When faced by contradictory conflicts between several, e.g., two, pieces of advice, humans, furthermore, often employ the **tactic** of resolving the conflict definitively in favor of the advice that arises from the *source* that has greater authority or reliability (or, more generally, some other property).

By authority in this paper, we mean in the legal and organizational senses. For example, federal law takes precedence over state law; directives from the head of an organization take precedence over those from subordinates.

By reliability in this paper, we mean in the sense of accuracy or likelihood

[1] as in the terminology "knowledge assimilation", for example cf. the AAAI Spring Symposium of that title in March 1992 held at Stanford University, Stanford, California

[2] More precisely, the agent's theory can be shown to be consistent given some reasonable, often-met syntactic conditions on the premises. In [Grosof, 1993], we discuss these conditions that ensure consistency.

of correctness. For example, I believe the New York Times newspaper to be a more reliable source than the Village Voice newspaper, which in turn I beleve to be more reliable than the National Enquirer newspaper. In [Grosof, 1991] [Grosof, 1992] (especially section 2.5), we elaborate more on what we mean by reliability and authority, and give more examples.

In this paper, we show how to formalize the second above-mentioned tactic, in combination with the first: in our approach, we infer prioritization-type precedence [Lifschitz, 1985] [Grosof, 1991] [Grosof, 1992] (chapter 2) from currently available information about the partial order of relative authority and/or reliability. To do this, we employ a formalism that we have recently developed: Defeasible Axiomatized Policy (DAP) circumscription. DAP circumscription can represent, moreover, the *non*-monotonic inference of prioritization. I.e., the prioritization may itself be defeasible.

2 An Actual Example of Taking Advice in a Bureaucracy

In this section, we illustrate by example that bureaucracy is one domain where such taking of advice from (potentially) contradictory multiple sources, according to the above two tactics, is important and useful. In dealing with a bureaucracy: one has to take actions in a somewhat mysterious realm with its own rules. Different bureaucrats tell one different things; sometimes their advice conflicts. However, one also has beliefs about which bureaucrats have greater authority (e.g., organizational rank) or reliability (e.g., competence) than others.

Defeasible Authority, Reliability, Precedence:
In general, these beliefs about authority and reliability, and thus about prioritization-type precedence, may themselves be **defeasible**. E.g., promotions and demotions may change relative authority. E.g., the progressing history of predictive accuracy may change relative reliability. For simplicity's sake, in the example below, we consider, however, only monotonic inferences about precedence. In [Grosof, 1993], we detail a second example that indeed has *non*-monotonic inferences about precedence.

Consider the following true story. Bold type in the remainder of this section indicates the most essential parts of the reasoning in the story. By "essential" here, we mean with respect to the problem addressed in this paper.

One Saturday a few months ago, one of the fluorescent overhead lights in my office at IBM suddenly stopped working. I knew that I had to make some kind of official request for the repair, but I did not know exactly who and how to ask. I thus phoned Security and asked them what was the procedure to officially request the lighting repair. The woman answering the phone told me (**Security advice**) that the procedure was simply to type the command "MAINT" to the local mainframe computer system, then obey the instructions it gave me. **I now believed that this was the correct procedure**; in general, in the past, I had found the Security people to be pretty reliable (i.e., competent). I thus acted by following her advice and did the MAINT command: this turned out to be merely to fill out a structured e-mail message to a relevant IBM bureaucrat in charge of maintenance-type repairs. After I sent the message, the system replied automatically that my request was being forwarded first to my departmental Administrative Assistant (AA), who had to approve the request.

Monday afternoon came, and I was just wondering impatiently about my lighting repair, when I received an e-mail message from my (departmental) Administrative Assistant. In it, she said that she had *not* approved my repair request, that it was inappropriate. She said (**AA advice**) that the correct procedure for requesting the lighting repair was to phone Security and to ask them to officially request on my behalf to the building landlord (a separate, non-IBM company, named Robert Martin), to have their maintenance people come and fix it. ("Classic!", I mumbled to myself.) **I now believed that this was the correct procedure. Though it contradicted (i.e., conflicted with) what the weekend Security person had told me, I believed that my AA had more organizational authority than Security on this kind of matter and, in the past, had generally been even more reliable (i.e., competent) than Security. Hence I resolved the conflict definitively (for the time being) in the AA's favor (more precisely, in favor of her advice.) That is, I assigned greater prioritization-type precedence to the AA advice than to the Security advice.**

I thus acted by following the AA's advice. (Though Security again repeated their previous advice, I insisted that they contact the landlord.) Sure enough, within fifteen minutes a landlord maintenance person came to my office and completed the repair.

3 DAP Circumscription

DAP circumscription is fairly complex to define. Next, we sketch it briefly. In [Grosof, 1993], we describe Defeasible Axiomatized Policy (DAP) circumscription in technical and intuitive detail.

By a default premise belief, we mean an open (first-order) formula that is maximized[3]; in Reiter's Default Logic [1980], this corresponds roughly to the idea of an open normal default without pre-requisite. In DAP circumscription, any default premise belief can be expressed via the minimization of a predicate (view it as an abnormality cf. [McCarthy, 1986] [Grosof, 1984] [Lifschitz, 1984]), or, more generally, via the minimization of a predicate at some subset of that predicate's set of "points", i.e., instances. Here, we mean point in the sense used by Lifschitz in his idea of pointwise circumscription [1987] [1988]. A special first-order predicate $N1$ specifies which points (of which predicates) are minimized. A special first-order predicate $R1$ specifies which minimized points have higher prioritization than which other minimized points. So as to define the activity of $N1$ and $R1$ appropriately, we introduce into the first-order ("base") language $\mathcal{L}$ the name of each predicate in $\mathcal{L}$. I.e., we introduce the new object symbol $'P$ for each predicate symbol P. This introduction of new symbols is essentially an extremely limited form of quotation.

As is usual in circumscription, a sequence of informational updates to an agent (here, the advice-taker) is represented as the accumulation of (a conjunctive collection of) *base axioms*. Each base axiom is a for-sure premise belief, expressed in the first-order base language $\mathcal{L}$. We call the conjunction of the base axioms: *the base sentence*. At each point in the sequence of updates, the current base sentence specifies a corresponding DAP circumscription, which represents the current non-monotonic theory. The DAP circumscription is a formula in second-order logic. At each point in the sequence, the agent's non-monotonic conclusions are exactly those which are entailed (according to second-order classical logic's model theory) by that DAP circumscription formula.

4 Formalization of the Example

In DAP circumscription, the lighting-repair example from section 2 can be formalized as follows. The advice-taking agent in this example is me.

Lower-case arguments indicate variables. sit stands for situation. Sec stands for Security. $MyPb$ stands for my particular problem situation. $Light$ stands for lighting repair. $todo(subj, sit, proced)$ stands for the proposition that the (unique) correct bureaucratic procedure is $proced$ when the subject is $subj$ and the situation is sit. $MAINT$ stands for the procedure of performing

[3] equivalently, a formula whose negation is minimized in the sense of what [McCarthy, 1986] called formula circumscription

the MAINT command on the mainframe. $Sec\_Ask\_Landlord$ stands for the procedure of asking Security to ask the landlord. $competent$ stands for basically competent and co-operative. $more\_reliable(source1, source2, subject)$ stands for $source1$ being more reliable than $source2$ about $subject$; similarly for $more\_authoritative$. $\approx\!\!\!|\ p$ indicates that the sentence p is non-monotonically concluded from the DAP circumscription specified by the current base sentence. $\exists ! x. \ldots$ stands for "there exists a *unique* x such that ...".

ab is a single overall abnormality predicate; it takes

$\langle source, subject, situation \rangle$

as (the type of) its argument tuple.

$N1('ab, source23, subject45, situation67)$

thus expresses the minimization of the point

$ab(source23, subject45, situation67)$.

$R1('ab, source1, subj1, sit1, 'ab, source2, subj2, sit2)$

thus expresses that the minimization of

$ab(source1, subj1, sit1)$

has higher prioritization than the minimization of

$ab(source2, subj2, sit2)$.

Before receiving the Security advice or the AA advice, the base sentence of the circumscription includes exactly (the conjunction of) the "background" base axioms: the uniqueness of names, plus

$\forall subject, sit. \exists ! procedure. todo(subject, sit, procedure)$

$\forall source, subject, sit. competent(source) \supset N1('ab, source, subject, sit)$

$competent(Sec)$

$competent(AA)$

$$\begin{aligned}
&\forall source1, source2, subject, sit. \\
&\qquad more_reliable(source1, source2, subject) \\
&\qquad \wedge\ more_authoritative(source1, source2, subject) \\
&\qquad\qquad \supset \\
&\qquad R1('ab, source1, subject, sit, 'ab, source2, subject, sit)
\end{aligned}$$

$more\_reliable(AA, Sec, Light)$

$more\_authoritative(AA, Sec, Light)$

These last three axioms represent the information from which I (the advice-taker) am able to conclude (monotonically) that the AA advice has higher prioritization than (i.e., takes precedence over) the Security advice. (As a practical matter, this conclusion is drawn only after the AA advice update below.)

The Security advice update corresponds to adding the base axiom:

$\forall sit. \neg ab(Sec, Light, sit) \equiv todo(Light, sit, MAINT)$

After this update, I non-monotonically conclude that I should do the MAINT command on the mainframe:

$\approx\ \ todo(Light, MyPb, MAINT)$

The later, AA advice update corresponds to adding the base axiom:

$\forall sit.\ \neg ab(AA, Light, sit) \ \equiv\ todo(Light, sit, Sec\_Ask\_Landlord)$

After this update, I retract the previous non-monotonic conclusion above, and non-monotonically conclude instead that I should have Security contact the landlord:

$\approx\ \ todo(Light, MyPb, Sec\_Ask\_Landlord)$

5 Method that the Example Illustrates

Next, we sketch a general method of formalizing advice-taking which the example in the last section illustrates. The method goes beyond just this example.

In the advice-taker / learner / assimilator:

1. Represent advice to be the content of utterance information from a useful source. A source which is basically honest and competent is useful.

2. Represent each piece of advice as a default schema, i.e., a default premise belief.

3. Represent this default schema via an abnormality predicate.

4. Represent a source as an index or index tuple for the default schema.

5. Represent precedence explicitly, so that one can reason about it, e.g., to infer it from reliability and authority information.

In [Grosof, 1993], we detail another example where precedence is defeasible, and formalize it in DAP circumscription in the above manner. The capabilities of DAP circumscription to represent "pointwise" detail about precedence, also discussed in [Grosof, 1993], are useful, as well. (By "pointwise", we mean in the sense introduced by Lifschitz [1987] [1988].)

6 Discussion, Related Work, and Future Work

The problem we address in this paper has not received much previous attention in the AI literature, especially in connection with common-sense reasoning. We are not aware of any previous work about advice-taking, or even common-sense reasoning, in the domain of bureaucracy.

There is a body of previous work on formalizing speech acts, between multiple agents, using non-monotonic logic, e.g., [Perrault, 1987] [Appelt and Konolige, 1988]. However, this work does not deal with the issue of conflict between the (content of the) proferred "pieces of advice" from multiple agents, nor with precedence between the sources of advice. In short, this body of work is not highly relevant.

Morgenstern and Guerriro [1991] address non-monotonic reasoning about other agents' non-monotonic reasoning and beliefs. Their concerns are quite different from ours: their focus is on each agent's beliefs about what other agents believe, not on what one agent should believe for itself when told things by the other agents. Our problem is simpler than theirs in one sense: we care only about how to use what the other agents tell the advice-taker agent, not what those other agents believe. Morgenstern and Guerriro do not deal with the (general) issue of conflict between the (content of the) proferred "pieces of advice" from multiple agents, nor with precedence between the sources of advice. In short, this body of work is not highly relevant, but it is complementary to ours.

By far the most relevant previous work is by Lifschitz [1988]. Lifschitz was the first and (until us) only to address prioritization-type precedence among multiple sources of default information. He developed a new variant of circumscription and showed how to monotonically infer precedence, and gave some examples of reliability as a basis for such reasoning about precedence. His work inspired our development of DAP circumscription. He did not focus mainly on the issue of advice-taking, however.

Compared to Lifschitz', our approach has several representational advantages. Firstly and most fundamentally, DAP circumscription is capable of expressing defeasible prioritization; Lifschitz' formalism is not.

Secondly, Lifschitz' formalism cannot directly represent prioritization that is not layered (i.e., not "stratified"). [Grosof, 1991] (elaborated in [Grosof, 1992]) showed that non-layered prioritization is very important for applications. While it is possible (using reduction theorems given in [Grosof, 1992]) to represent non-layered prioritization in Lifschitz' formalism in an indirect manner, it is often rather awkward, involving a substantial increase in the expressive complexity (e.g., clausality / universality may be lost) and size (e.g., squared blow-up) of the representation.

Thirdly, our approach also represents the concepts of defaults and priorities directly, unlike Lifschitz' which uses only the less intuitive concept of varying versus not varying. Ours is more epistemologically natural and thus facilitates knowledge acquisition.

Fourthly, Lifschitz' formalism permits "nasty" unsatisfiability of an extra kind. By nasty unsatisfiability, we mean the situation in which the circumscription is unsatisfiable, even though the base sentence is satisfiable; i.e., the circumscription operation itself introduces inconsistency. Circumscription and other model-preference NM formalisms are well-known to be susceptible to nasty unsatisfiability: when models are not "well-founded" in the sense defined by [Etherington *et al.*, 1985] [Lifschitz, 1986]. But Lifschitz' formalism permits a kind of incoherence, in the "policy" of varying, which results in nasty unsatisfiability even in very simple, "well-founded" theories. Our approach does not.

Fifthly, because our approach, unlike Lifschitz', is a special case of single general circumscription [Lifschitz, 1984], it is easier to apply results available from the body of work on model-preference logics (e.g., [Shoham, 1988] [Brown and Shoham, 1989] [Boddy *et al.*, 1989] [Selman and Kautz, 1989]), and on preferential logics (e.g., [Kraus *et al.*, 1990] [Lehmann, 1989])

See [Grosof, 1992] (section 8.3) and [Grosof, 1993] for more detailed review of Lifschitz' approach and comparison of our approach to his.

Besides Lifschitz', other previous non-monotonic logical formalisms cannot represent precedence explicitly, e.g., to infer it; or cannot represent it at all. [Grosof, 1993] elaborates on this contrast.

In terms of common-sense reasoning, this work fits into the agenda of Hobbs [1985], who says (page xv): "Much, much more needs to be done in commonsense psychology and sociology. ...The cornerstone of theories of the social world must be axiomatizations of mutual belief and the domain of communicative acts."

In terms of machine learning, this work is an interesting kind of learning with prior knowledge, where a powerful kind of induction can be justified and performed: generalizing universally from a single "example" instance. E.g., in sections 2 and 4, after taking the AA advice, the universal rule can be concluded (non-monotonically)that the correct procedure in *any* situation of lighting repair is to ask Security to contact the landlord.

Learning by taking advice is a kind of learning "in the small", by contrast with, e.g., empirical induction from a large number of instances.

We would argue that some kinds of learning are part of what are called "common-sense reasoning and abilities". Carbonell and Minton [1985] also argue for this point.

In terms of domains, bureaucracy is important and interesting in its own right. With respect to the issues raised in this paper, it is, moreover, representative of many domains involving offices.

One possible alternative approach to handling the potentially contradictory multiplicity of advice is to maintain more of an arm's length distance from advice to one's beliefs. That is, one might *not* adopt things one is told so directly as one's own beliefs. The difficulty with this approach is that one's beliefs may then not be strong enough to act upon. In this paper, we have in mind applications of advice-taking in which the need to decide between actions implies a pressure to commit: to decide which advice to adopt as actionable beliefs. One might take a decision-theoretic[4] approach, however, in which such commitment occurs only at decision time. Such an approach requires information that is more complex in some respects: probabilities and utilities. The costs and benefits of different pieces (or sources) of advice being right versus wrong for the particular decisions at hand can then be probabilistically weighed and balanced.

Pursuing such a decision-theoretic approach is an interesting direction for future work. However, such an approach cannot escape the issues raised in this paper. Consider: where does the information about the probabilities and utilities come from? We would argue that, typically, such information is itself defeasible and arises itself from a multiplicity of sources.

Another important direction for future work is to explore how to mechanize DAP circumscription. In [Grosof, 1993], we discuss this more. We have some results showing that DAP inference often (e.g., in our example formalization in section 4) reduces to inference in ordinary prioritized predicate circumscription cf. [Grosof, 1991].

[4]By decision-theoretic, we mean in the sense of maximizing probabilistically-expected utility, as pioneered in mathematical economics, e.g., cf. [Raiffa, 1968].

References

[Appelt and Konolige, 1988] Douglas Appelt and Kurt Konolige. A practical nonmonotonic theory for reasoning about speech acts. Technical Report Technical Note 432, SRI International, 333 Ravenswood Ave., Menlo Park, California 94025, Apr 1988.

[Boddy *et al.*, 1989] Mark Boddy, Robert P. Goldman, Keiji Kanazawa, and Lynn A. Stein. Investigations of model-preference defaults. Working paper, submitted to Computational Intelligence journal., 1989.

[Brown and Shoham, 1989] Allen L. Brown, Jr. and Yoav Shoham. New results on semantical nonmonotonic reasoning. In M. Reinfrank *et al*, editor, *Proceedings of the Second International Workshop on Non-Monotonic Reasoning*, Berlin, Germany, 1989. Springer Lecture Notes on Computer Science.

[Carbonell and Minton, 1985] Jaime G. Carbonell and Steven Minton. Metaphor and commonsense reasoning. In Jerry R. Hobbs and Robert C. Moore, editors, *Formal Theories of the Commonsense World*, pages 405–426. Ablex Publishing Corporation, 355 Chestnut Street, Norwood, NJ 07648, 1985.

[Dietterich, 1982] Thomas G. Dietterich. Learning and inductive inference. In Paul R. Cohen and Edward A. Feigenbaum, editors, *The Handbook of Artificial Intelligence, Volume 3*, pages 323–512. Morgan Kaufmann, San Mateo, California, 1982.

[Etherington *et al.*, 1985] D. Etherington, R. Mercer, and R. Reiter. On the adequacy of predicate circumscription for closed-world reasoning. *Computational Intelligence*, 1:11–15, 1985.

[Grosof, 1984] Benjamin N. Grosof. Default reasoning as circumscription. *Proceedings of the First AAAI Non-Monotonic Reasoning Workshop*, pages 115–124, Oct 1984. Held New Paltz, NY.

[Grosof, 1991] Benjamin N. Grosof. Generalizing prioritization. *Proceedings of the Second International Conference on Principle of Knowledge Representation and Reasoning*, pages 289–300, April 1991. Also available as IBM Research Report RC15605, IBM T.J. Watson Research Center, P.O. Box 704, Yorktown Heights, NY 10598.

[Grosof, 1992] Benjamin N. Grosof. *Updating and Structure in Non-Monotonic Theories*. PhD thesis, Computer Science Dept., Stanford University, Stanford, California 94305, October 1992. Published by University Microfilms, Inc.. Also available as Research Report from Stanford University Computer Science Dept. and/or IBM T.J. Watson Research Center.

[Grosof, 1993] Benjamin N. Grosof. Defeasible and pointwise prioritization: Preliminary report. Technical report, IBM T.J. Watson Research Center, P.O. Box 704, Yorktown Heights, NY 10598, January 1993. Revised from Working Paper of April 1992.

[Hobbs, 1985] Jerry R. Hobbs. Introduction. In Jerry R. Hobbs and Robert C. Moore, editors, *Formal Theories of the Commonsense World*, pages xi–xxii. Ablex Publishing Corporation, 355 Chestnut Street, Norwood, NJ 07648, 1985.

[Kraus *et al.*, 1990] S. Kraus, D. Lehmann, and M. Magidor. Preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.

[Lehmann, 1989] D. Lehmann. What does a conditional knowledge base entail? *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning*, pages 212–222, Toronto, Ontario, 1989. Morgan Kaufmann.

[Lifschitz, 1984] Vladimir Lifschitz. Some results on circumscription. *Proceedings of the First AAAI Non-Monotonic Reasoning Workshop*, pages 151–164, Oct 1984. Held New Paltz, NY.

[Lifschitz, 1985] V. Lifschitz. Computing circumscription. *Proceedings IJCAI-85*, pages 121–127, Los Angeles, CA, 1985.

[Lifschitz, 1986] V. Lifschitz. On the satisfiability of circumscription. *Artificial Intelligence*, 28:17–27, 1986.

[Lifschitz, 1987] Vladimir Lifschitz. Pointwise circumscription. In Matthew L. Ginsberg, editor, *Readings in Nonmonotonic Reasoning*. Morgan Kaufmann, San Francisco, CA, 1987.

[Lifschitz, 1988] V. Lifschitz. Circumscriptive theories: a logic-based framework for knowledge representation. *Journal of Philosophical Logic*, 17:391–441, 1988.

[McCarthy, 1968] John McCarthy. Programs with common sense. In Marvin Minsky, editor, *Semantic Information Processing*, pages 403–409, Cambridge, Mass., 1968. MIT Press. Reprinted from original 1958 appearance in: Proceedings of the Symposium on the Mechanization of Thought Processes, National Physical Laboratory I:77–84.

[McCarthy, 1980] J. McCarthy. Circumscription—a form of non-monotonic reasoning. *Artificial Intelligence*, 13:27–39, 1980.

[McCarthy, 1986] J. McCarthy. Applications of circumscription to formalizing commonsense knowledge. *Artificial Intelligence*, 28:89–116, 1986.

[Morgenstern and Guerriro, 1991] Leora Morgenstern and Ramiro Guerriro. Epistemic logics for multiple agent nonmonotonic reasoning I. Submitted. IBM T.J. Watson Research Center, P.O. Box 704, Yorktown Heights, NY 10598. Includes revised version of contents of paper appearing in Proceedings of AAAI-90, entitled "A Formal Theory of Multiple Agent Nonmonotonic Reasoning", by Leora Morgenstern., Sep 1991.

[Perrault, 1987] C. Raymond Perrault. An application of Default Logic to speech act theory. Technical Report CSLI-87-90, Center for the Study of Language and Information, Stanford University, Ventura Hall, Stanford University, Stanford, California 94305, Mar 1987.

[Raiffa, 1968] Howard Raiffa. *Decision Analysis: Introductory Lectures on Choices under Uncertainty.* Addison-Wesley, Reading, Mass., 1968.

[Reiter, 1980] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 12:81–132, 1980.

[Selman and Kautz, 1989] Bart Selman and Henry A. Kautz. Model-preference default theories. Technical Report KRR-TR-89-7, Dept. of Computer Science, Univ. of Toronto, Toronto, Ontario, Canada M5S 1A4, Apr 1989.

[Shoham, 1988] Y. Shoham. *Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence.* MIT Press, Cambridge, Mass., 1988.