BENCHMARK PROBLEMS
FOR FORMAL NONMONOTONIC REASONING

Version 2.00

Vladimir Lifschitz

Stanford University
Stanford, CA 94305, USA

Research on the theory of nonmonotonic reasoning has given us several important formalisms and many valuable ideas on the methodology of their use. Much work has been done on the investigation of the possibilities and limitations of different approaches. An important role in this work belongs to some "benchmark" examples of nonmonotonic commonsense reasoning. This paper contains a long list of problems of this kind. Hopefully, the evaluation of future progress can be facilitated with a list like this available for reference.

The description of each problem consists of a list of *assumptions*, followed by one or more *conclusions*. The goal is to represent each of the assumptions and conclusions by an expression in a formal language with a declarative semantics, and to verify that each conclusion is indeed a consequence of the assumptions. (In most cases, there is a general consensus in the nonmonotonic community about the validity of the patterns of reasoning exemplified in these problems, but there are exceptions.) Besides the assumptions explicitly included in the statement of the problem, some other commonsense facts, "implicit assumptions," may have to be used. A discussion of the "rules of the game" accepted in the theory of nonmonotonic reasoning can be found in Section 1.3.1 of (Ginsberg 1987).

It should be emphasized that the problems illustrate conceptual, rather than computational difficulties. In each example, the challenge is to formalize it—not implement it on the computer.

John McCarthy recommends that the following goals be kept in mind when nonmonotonic systems are used for formalizing commonsense reasoning (personal communication, June 1988):

Generality. The formal expression of general facts should be suitable for inclusion in a general commonsense database. They should not be ad hoc to the particular example.

Elaboration tolerance. When additional facts are added that shouldn't affect the conclusion, then the conclusion should not be affected. Ideally, the length of the reasoning process shouldn't grow much.

Locality of reasoning. As much of the reasoning as possible should involve small numbers of facts, even though in general nonmonotonic reasoning requires taking into account the whole set of facts.

The list of problems is followed by an appendix, containing several solutions to the first problem on the list. Some solutions use the same nonmonotonic system, but in different ways. We haven't tried to illustrate all important formalisms described in the literature; the number of solutions could be easily doubled. But even the solutions that are included show the considerable variety of languages and methodologies available now. Why do we need so many different approaches? The main reason is that some approaches work better than others when applied to more complex problems. Ideally, we would like to have a single system of nonmonotonic reasoning that leads to correct and concise solutions to all benchmark examples. But formalisms with limited possibilities can be valuable too, if they are computationally tractable or particularly easy to use, and if their relation to more expressive systems is well understood.

The notes at the end of the paper contain brief comments on the particular difficulty emphasized in each example, and give selected references to the literature.

To be useful, this list of benchmark problems should grow and change, as researchers turn to more complex forms of commonsense reasoning and to finer distinctions between possible meanings of informally stated examples. Any comments, criticisms and contributions will be greatly appreciated.

A. Default Reasoning

A1. Basic Default Reasoning.

Assumptions: Blocks A and B are heavy.

Heavy blocks are normally located on the table.

A is not on the table.

Conclusion: B is on the table.

A2. Default Reasoning with Irrelevant Information.

Assumptions: Blocks A and B are heavy.

Heavy blocks are normally located on the table.

A is not on the table.

B is red.

Conclusion: B is on the table.

A3. Default Reasoning with Several Defaults.

Assumptions: Blocks A and B are heavy.

Heavy blocks are normally located on the table.

Heavy blocks are normally red.

A is not on the table.

B is not red.

Conclusions: B is on the table.

A is red.

A4. Default Reasoning with a Disabled Default.

Assumptions: Blocks A and B are heavy.

Heavy blocks are normally located on the table.

A is possibly an exception to this rule.

Conclusion: B is on the table.

A5. Default Reasoning in an Open Domain.

Assumptions: Block A is heavy.

Heavy blocks are normally located on the table.

A is not on the table.

Conclusion: All heavy blocks other than A are on the table.

A6. Reasoning about Unknown Exceptions I.

Assumptions: Blocks A, B and C are heavy.
Heavy blocks are normally located on the table.
At least one of A, B is not on the table.

Conclusions: C is on the table.
Exactly one of A, B is not on the table.

A7. Reasoning about Unknown Exceptions II.

Assumptions: Heavy blocks are normally located on the table.
At least one heavy block is not on the table.

Conclusion: Exactly one heavy block is not on the table.

A8. Reasoning about Unknown Exceptions III.

Assumptions: Block A is heavy.
Heavy blocks are normally located on the table.
At least one heavy block is not on the table.

Conclusion: A is on the table.

A9. Priorities between Defaults.

Assumptions: Jack asserts that block A is on the table.
Mary asserts that block A is not on the table.
When Jack asserts something, he is normally right.
When Mary asserts something, she is normally right.
Mary's evidence is more reliable that Jack's.

Conclusion: Block A is not on the table.

A10. Priorities between Instances of a Default.

Assumptions: Jack asserts that block A is on the table.
Mary asserts that block A is not on the table.
When people assert something, they are normally right.
Mary's evidence is more reliable that Jack's.

Conclusion: Block A is not on the table.

A11. Reasoning about Priorities.

Assumptions: Jack asserts that block A is on the table.
Mary asserts that block A is not on the table.
When people assert something, they are normally right.

Conclusion: If Mary's evidence is more reliable that Jack's, then block A is not on the table.

B. Inheritance

B1. Linear Inheritance.

Assumptions: Animals normally do not fly.
Birds are animals.
Birds normally fly.
Ostriches are birds.
Ostriches normally do not fly.

Conclusions: Animals other than birds do not fly.
Birds other than ostriches fly.
Ostriches do not fly.

B2. Tree-Structured Inheritance.

Assumptions: Animals normally do not fly.
Birds are animals.
Birds normally fly.
Bats are animals.
Bats normally fly.
Ostriches are birds.
Ostriches normally do not fly.

Conclusions: Animals other than birds and bats do not fly.
Birds other than ostriches fly.
Bats fly.
Ostriches do not fly.

B3. One-Step Multiple Inheritance.

Assumptions: Quakers are normally pacifists.
Republicans are normally not pacifists.

Conclusions: Quakers who are not Republicans are pacifists.
Republicans who are not Quakers are not pacifists.

B4. Multiple Inheritance.

Assumptions: Quakers are normally pacifists.
Republicans are normally hawks.
Pacifists are normally politically active.
Hawks are normally politically active.
Pacifists are not hawks.

Conclusions: Quakers who are not Republicans are pacifists.
Republicans who are not Quakers are hawks.
Quakers, Republicans, pacifists and hawks are politically active.

C. Uniqueness of Names

C1. Unique Names Hypothesis for Objects.

Assumptions: Different names normally denote different objects.

The names "Ray" and "Reiter" denote the same person.

The names "Drew" and "McDermott" denote the same person.

Conclusion: The names "Ray" and "Drew" denote different people.

C2. Unique Names Hypothesis for Functions.

Assumptions: Different people normally have different fathers.

Joseph and Benjamin have the same father.

Gaius and Tiberius have the same father.

Conclusion: Joseph and Gaius have different fathers.

D. Reasoning about Action

D1. Frame Problem for Temporal Projection.

Assumptions: After an action is performed, things normally remain as they were.

Any time the robot grasps a block, the block will be in the hand.

If a block is in the hand, then, after the robot moves it onto the table, the block will be on the table.

Initially block A is not in the hand.

Initially block A is not on the table.

Conclusion: After the robot grasps block A, waits, and then moves it onto the table, the block will be on the table.

D2. Temporal Projection.

Assumptions: After an action is performed, things normally remain as they were.

When the robot grasps a block, the block will be normally in the hand.

When the robot moves a block onto the table, the block will be normally on the table.

Moving a block that is not in the hand is an exception to this rule.

Initially block A is not in the hand.

Initially block A is not on the table.

Conclusion: After the robot grasps block A, waits, and then moves it onto the table, the block will be on the table.

D3. Temporal Projection with Ramifications.

Assumptions: After an action is performed, things normally remain as they were.

A block is on the table if and only if it is not on the floor.

When the robot grasps a block, the block will be normally in the hand.

When the robot moves a block onto the table, the block will be normally on the table.

Moving a block that is not in the hand is an exception to this rule.

Initially block A is not in the hand.

Initially block A is on the floor.

Conclusion: After the robot grasps block A, waits, and then moves it onto the table, the block will not be on the floor.

D4. Temporal Explanation with Unknown Initial Conditions.

Assumptions: After an action is performed, things normally remain as they were.

When the robot grasps a block, the block will be normally in the hand.

When the robot moves a block onto the table, the block will be normally on the table.

Moving a block that is not in the hand is an exception to this rule.

Initially block A was not on the table.

After the robot moved A onto the table and then waited, A was on the table.

Conclusion: Initially A was in the hand.

D5. Temporal Explanation with Unknown Actions.

Assumptions: After an action is performed, things normally remain as they were.

When the robot grasps a block, the block will be normally in the hand.

When the robot moves a block onto the table, the block will be normally on the table.

Moving a block that is not in the hand is an exception to this rule.

Initially block A was not on the table.

Initially block A was not in the hand.

After the robot grasped some block and then moved some block onto the table, A was on the table.

Conclusions: The block that was grasped was A.

The block that was moved onto the table was A.

D6. Temporal Explanation with Actions of Unknown Kinds.

Assumptions: After an action is performed, things normally remain as they were.

When the robot grasps a block, the block will be normally in the hand.

When the robot moves a block onto the table, the block will be normally on the table.

Moving a block that is not in the hand is an exception to this rule.

Initially block A was not on the table.

Initially block A was not in the hand.

After the robot performed two actions, A was on the table.

Conclusions: The first of the two actions was grasping A.

The second of the two actions was moving A onto the table.

D7. Reasoning about the Unknown Order of Actions.

Assumptions: After an action is performed, things normally remain as they were.

When a person accepts a job offer from some employer, he will be employed by that employer.

If Bill is offered a job at Berkeley or at Stanford when he is unemployed, he will accept it.

Bill is currently unemployed.

Conclusion: After Bill is offered jobs at Berkeley and Stanford at two different instants of time, he will be employed either by Berkeley or by Stanford.

D8. Reasoning about Unexpected Change.

Assumptions: After an action is performed, things normally remain as they were.

When the robot moves a block to another location, the block will be normally at that location.

After the robot moved a block to Location 1 and then to Location 2, the block changed its color.

Conclusion: The block changed its color only once, either after the first move, or after the second.

D9. Reasoning about the Unexpected Absence of Change.

Assumptions: When the robot moves a block to another location, the block will be normally at that location.

After the robot moved block A onto the table, and then moved block B onto the table, at most one of the blocks A, B was on the table.

Conclusion: After the two actions were performed, exactly one of the blocks A, B was on the table.

D10. Counterfactual Reasoning about Unexpected Change.

Assumptions: After an action is performed, things normally remain as they were.

When the robot moves a block to another location, the block will be normally at that location.

After the robot moved block A to Location 1, block B changed its color.

Conclusion: Block B would have changed its color if the robot had moved A to Location 2.

D11. Reasoning about Concurrent Actions.

Assumptions: When two actions are performed concurrently, their effects are normally combined.

After a block is moved to another location, it is normally at that location.

Conclusion: After block A is moved to Location 1, and block B is concurrently moved to Location 2, A will be at Location 1 and B will be at Location 2.

E. Autoepistemic Reasoning

E1. Basic Autoepistemic Reasoning.

Assumption: Block A is on the table.

Conclusion: It is not known whether B is on the table.

E2. Autoepistemic Reasoning with Incomplete Information.

Assumption: At least one of the blocks A, B is on the table.

Conclusions: It is not known whether A is on the table.

It is not known whether B is on the table.

E3. Autoepistemic Reasoning in an Open Domain.

Assumption: Block A is on the table.

Conclusion: About any block other than A it is not known whether it is on the table.

E4. Autoepistemic Default Reasoning.

Assumptions: Blocks that are not known to be heavy are on the table.

Block A is heavy.

Conclusion: Block B is on the table.

Appendix. Seventeen Solutions to Problem A1

The first four solutions are based on default logic (Reiter 1980).

Solution 1. The default theory with the axioms

$$heavy\ A, \quad heavy\ B, \tag{1}$$

$$\neg ontable\ A \tag{2}$$

and the default

$$\frac{heavy\ x\ :\ \mathrm{M}\ ontable\ x}{ontable\ x}. \tag{3}$$

Instead of (3), a default without a prerequisite can be used:

Solution 2. The default theory with axioms (1) and (2), and with the default

$$\frac{:\ \mathrm{M}\ heavy\ x \supset ontable\ x}{heavy\ x \supset ontable\ x}. \tag{4}$$

The difference between these two methods is discussed in (Imielinski 1987), Section 3.2.

Default logic can be also used in combination with abnormality predicates proposed in (McCarthy 1986):

Solution 3. The default theory with axioms (1), (2) and

$$heavy\ x \wedge \neg ab\ x \supset ontable\ x, \tag{5}$$

and with the default

$$\frac{:\ \mathrm{M}\ \neg ab\ x}{\neg ab\ x}. \tag{6}$$

There is also a quite different approach, based on nonnormal defaults:

Solution 4. The default theory with axioms (1), (2) and

$$ab\ A, \tag{7}$$

and with the default

$$\frac{heavy\ x\ :\mathrm{M}\ \neg ab\ x}{ontable\ x}. \tag{8}$$

The usefulness of nonnormal defaults for some applications was noted by Morris (1987).

Next we will consider some counterparts of Solutions 3 and 4 in autoepistemic logic (Moore 1985). Since Moore's system is propositional, all formulas with variables have to be replaced by their ground instances. For instance, (5) becomes

$$heavy\ \xi \wedge \neg ab\ \xi \supset ontable\ \xi \qquad (\xi \in \{A, B\}). \tag{9}$$

Solution 5. The autoepistemic theory with axioms (1), (2), (9) and

$$ab\ \xi \supset \mathrm{L}\ ab\ \xi \qquad (\xi \in \{A, B\}). \tag{10}$$

Solution 6. The autoepistemic theory with axioms (1), (2), (7) and

$$heavy\ \xi \wedge \neg \mathrm{L}\ ab\ \xi \supset ontable\ \xi \qquad (\xi \in \{A, B\}). \tag{11}$$

The method used in the last solution was proposed by Gelfond (1988).

Konolige (1987) proposed a general method for translating default theories into autoepistemic theories. For instance, when applied to Solution 4, this transformation gives the following modification of Solution 6:

Solution 7. The autoepistemic theory with axioms (1), (2), (7) and

$$\mathrm{L}\ heavy\ \xi \wedge \neg \mathrm{L}\ ab\ \xi \supset ontable\ \xi \qquad (\xi \in \{A, B\}).$$

(Generally, Konolige's translation requires that quantifiers be added to autoepistemic logic, but in this case quantifiers are not needed.)

In Solutions 4, 6 and 7 the "cancellation of inheritance" axiom (7) is necessary in order to ensure the existence of a stable expansion—the difficulty that plagues the formalizations based on nonnormal defaults and on corresponding autoepistemic theories. This is discussed by Morris (1988), who proposed "stable closures" as an alternative to stable expansions as the basic concept of autoepistemic logic. The use of stable closures allows us to do without axiom (7), and we can simplify Solution 6 as follows:

Solution 8. The autoepistemic theory with axioms (1), (2) and (11).

The next solution is based on circumscription (McCarthy 1986). Here we need the "uniqueness of names" axiom

$$A \neq B. \tag{12}$$

Without it, we would be able to prove $\forall x(x \neq A \supset ontable\ x)$, but not *ontable* B.

Solution 9. The circumscriptive theory with axioms (1), (2), (5) and (12), with the predicate *ab* circumscribed, and with the predicate *ontable* varied.

In (Lifschitz 1988), a modification of circumscription is proposed that describes axiomatically which predicates are circumscribed or varied. Here is a solution based on that formalism:

Solution 10. The circumscriptive theory with axioms (1), (2), (5), (12) and

$$V[ab : ab, ontable].$$

The closed-world assumption (Reiter 1978), which models nonmonotonicity in deductive databases, seems to be too special even for our very first problem A1. But we can use the generalization of this concept proposed in (Genesereth and Nilsson 1987), Section 6.1—the closed-world assumption with respect to a predicate:

Solution 11. The theory with axioms (1), (2) and (5), and with the closed-world assumption applied to *ab*.

The predicate completion method (Clark 1978) requires that each given fact be represented by a clause with a distinguished positive literal; the relation of this literal is the relation that the clause is "about." Predicate completion can be viewed as a declarative semantics for logic programs with negation. In logic programming, axioms are called "rules," and each axiom containing more than one literal is written as the conditional whose consequent ("head") is its distinguished positive literal. For instance, the syntactic form of (5) shows that it is viewed as a fact about *ontable*.

Solution 12. The completion of the logic program (1), (5), (7).

Many other declarative approaches to the semantics of programs with negation can be used in this example instead of predicate completion. The important "iterated fixed point" semantics (Apt *et al.* 1988) is applicable only to the subclass of "stratified" programs, and the program in question is indeed stratified:

Solution 13. The stratified logic program with rules (1), (5) and (7).

Instead of logic programming languages, some frame languages can be used. The following definitions are written in the frame language from (Brewka 1987):

Solution 14. The definitions

(**defframe** Heavyblock (**slots** (Ontable True))),
(**definstance** A **of** Heavyblock **with** Ontable = False),
(**definstance** B **of** Heavyblock).

The next solution uses the assumption-based truth maintenance system from (de Kleer 1986). A declarative semantics for ATMSs is given in (Reiter and de Kleer 1987).

Solution 15. Facts (1) and (2), and the assumptions

$$heavy\ \xi \supset ontable\ \xi \qquad (\xi \in \{A, B\}).$$

A solution based on the language THEORIST (Poole 1988):

Solution 16. The THEORIST program

default heavy-blocks-are-on-table(X) : ontable(X) $\leftarrow$ heavy(X).
fact heavy(a).
fact heavy(b).
fact $\neg$ontable(a).

The last solution uses the multivalued logic approach (Ginsberg 1988).

Solution 17. The truth function assigning t to (1) and (2) and dt to

$$heavy\ x \supset ontable\ x$$

for all x.

Notes

Re A2. In comparison with A1, we have one additional assumption, and the conclusion remains the same. Nonmonotonicity means that an additional assumption may destroy a conclusion; but this shouldn't happens too easily, or else the formalism will be unusable.

Re A3. This example illustrates that there may be different "kinds" of abnormality, so that it would be a mistake, for instance, to express the assumptions about the typical location and color of a heavy block by (5) and

$$heavy\ x \wedge \neg ab\ x \supset red\ x.$$

We have to use "aspects," as proposed in (McCarthy 1986), or we can introduce two different abnormality predicates.

Re A4. Instead of assuming that A is on the table, as in A1, we want to make it impossible to determine where A is located. This difference can be expressed, for instance, by substituting (7) for (2) in Solution 3 or Solution 9.

Re A5. Emphasis here is on getting a universally quantified conclusion

$$\forall x(heavy\ x \wedge x \neq A \supset ontable\ x),$$

and not merely

$$heavy\ \xi \supset ontable\ \xi \qquad (\xi \neq A).$$

The difference is essential in the absence of a "domain closure assumption" like

$$\forall x(block\ x \equiv x = A \vee x = B \vee x = C)$$

(expressing that a complete list of blocks is available). This example presents a problem for "propositional" approaches, and also for Reiter's logic of defaults (because it treats an open default as the set of its closed instances; compare (McCarthy 1980), Remark 2). We should add that some researchers do not accept this form of reasoning as valid: If there is one exception to a default, then perhaps we can expect more (Matthew Ginsberg, personal communication, June 1988). This leads to the problem of formalizing "normally" in such a way that the conclusion of A1 would follow, but the conclusion of A5 wouldn't.

Re A6. Formalizations based on the closed-world assumption, similar to Solution 10 above, can become inconsistent in the presence of "disjunctive information" of this sort (Reiter 1978). This has led to various extensions of the original closed world assumption; for references, see (Ross and Topor 1987), Section 3.5. Of interest may be also a weaker interpretation of "normally," which permits inferring the first conclusion, but not the second; compare (Reiter 1984) and (Ross and Topor 1987).

Re A7. As in A5, it is essential here that quantifiers are needed for expressing the conclusion.

Re A8. Perlis (1986) pointed out that this is a difficult question. Even if we established, as in A7, that there is only one exception, how do we know that it is not A?

Re A9. This is not really "reasoning about belief"; the third and fourth assumptions can be expressed by something like

$$\textit{ontable-according-to-Jack}\ x \land \neg ab1\ x \supset \textit{ontable}\ x$$

and

$$\textit{ontable-according-to-Mary}\ x \land \neg ab2\ x \supset \textit{ontable}\ x.$$

Then the last assumption can be expressed by using prioritized circumscription (McCarthy 1986), (Lifschitz 1985). But this is perhaps not quite "declarative," and for this reason we may prefer an approach along the lines of Solution 10 (see (Lifschitz 1988), Section 12).

Re A10. This differs from A9 in that people should be "reified", and the axioms given in the previous note should be replaced by a single axiom like

$$\textit{ontable-according-to}(x, y) \land \neg ab(x, y) \supset \textit{ontable}\ x.$$

Then prioritized circumscription as defined in (Lifschitz 1985) won't be sufficient. See (Lifschitz 1988), Section 19.

Re A11. In this example, the importance of expressing priorities declaratively is especially evident. See (Lifschitz 1988), Section 19.

Re B1 and B2. See (Etherington and Reiter 1983); (McCarthy 1986), Sections 5 and 12. Most papers that study the semantics of inheritance hierarchies without relating them to general theories of nonmonotonic reasoning do not allow mixing strict inheritance ("birds are animals") and defeasible inheritance ("birds normally fly"), and consequently cannot handle these "heterogeneous" examples. (Horty and Thomason 1988) is an exception.

Re B3. This is the famous "Nixon diamond" (due to Raymond Reiter)—without Nixon. See (McCarthy 1986), Section 7.

Re B4. This enhancement of the Nixon diamond is due to Matthew Ginsberg.

Re C1 and C2. Although the unique names assumption is essentially the closed-world assumption applied to equality, its formalization is a difficult problem. See (Etherington *et al.* 1985), Section 5, and (McCarthy 1986), Section 6.

Re D1. This is a nonviolent version of the "Yale shooting problem" (Hanks and McDermott 1986). It can be assumed for simplicity that A is the only available block, so that the actions performed by the robot can be represented simply by constants *grasp*, *wait* and *move*. References to the extensive literature on this problem, with a critical discussion of the proposed solutions, can be found in (Hanks and McDermott 1987); see also (Morris 1987) and (Gelfond 1988).

Re D2. This enhancement of the previous example includes the "qualification problem," that motivated some of the early work on formal nonmonotonic reasoning (McCarthy 1980). The method of (Lifschitz 1987) addresses both the frame problem and the qualification problem.

Re D3. The assumption that describes the effect of moving a block onto the table is expressed in terms of one of two interrelated properies—being on the table and being on the floor. The fact that the other property is affected too is a "ramification" of this assumption. This is a simple example of the important and difficult "ramification problem," pointed out to me by Matthew Ginsberg.

Re D4, D5 and D6. Temporal projection is the simplest form of reasoning about action, important in view of its connection with planning; reasoning about the past, or "temporal explanation," is another possibility.

Re D7. Problems like this are sometimes said to involve "partially ordered events"; but what is partial here is, strictly speaking, the *available information* about the temporal order of events. This particular example is a product of my conversation with Murray Shanahan in August of 1988.

Re D8 and D9. See (Morgenstern and Stein 1988), (Lifschitz and Rabinov 1988).

Re D10. This is based on some examples of Matthew Ginsberg.

Re D11. This is based on examples of John McCarthy and Michael Gelfond. The first assumption is stated as a default, because there can be exceptions—for instance, trying to move the same object to two different places at the same time. At present, there seems to be no published completely formal work on nonmonotonic reasoning about concurrent actions.

Re E1 and E2. According to (Moore 1985), examples like these are quite different from the examples given before, in that nonmonotonicity stems here from a different source; these are examples of "autoepistemic," rather than "default" reasoning. Indeed, the word "known" is responsible for nonmonotonicity here, and not "normally," as before. But it is not clear how essential this distinction is. Some approaches to formalizing default reasoning use an "autoepistemic interpretation" of defaults; see Solutions 5–8 above. Moreover, some formal systems of default logic and autoepistemic logic turned out to be isomorphic (Konolige 1987).

Re E3. This requires predicate autoepistemic logic; see (Levesque 1988).

Re E4. This puzzling hybrid is suggested by a problem due to David Poole (personal communication, October 1987).

Acknowledgements

The idea of creating a list of benchmark problems for formal nonmonotonic reasoning came from John McCarthy. Matthew Ginsberg provided detailed comments on an earlier draft of this paper and suggested two of the solutions given in the appendix. Some problems included here have emerged in group discussions; it would be difficult to acknowledge the contributions of all individuals who took part in those discussions, and I will not even try. Compiling a list like this is necessarily a collective enterprise.

References

Apt, Krzysztof, Blair, Howard and Walker, Adrian (1988). Towards a theory of declarative knowledge, in: Minker, J. (ed.), *Foundations of Deductive Databases and Logic Programming*, Los Altos, CA: Morgan Kaufmann, pp. 89–148.

Brewka, Gerhard (1987). The Logic of Inheritance in Frame Systems, *Proc. IJCAI-87*, pp. 483–488.

Etherington, David; Mercer, Robert and Reiter, Raymond (1985). On the Adequacy of Predicate Circumscription for Closed-World Reasoning, *Computational Intelligence 1*, pp. 11–15. Reprinted in (Ginsberg 1987).

Etherington, David and Reiter, Raymond (1983). On Inheritance Hierarchies with Exceptions, *Proc. AAAI-83*, pp. 104–108. Reprinted in (Ginsberg 1987).

Clark, Keith (1978). Negation as Failure, in Gallaire, H. and Minker, J. (eds.), *Logic and Databases*, New York: Plenum Press, pp. 293–322. Reprinted in (Ginsberg 1987).

Gelfond, Michael (1988). Autoepistemic Logic and Formalization of Commonsense Reasoning, this volume.

Genesereth, Michael and Nilsson, Nils (1987). *Logical Foundations of Artificial Intelligence*, Los Altos, CA: Morgan Kaufmann.

Ginsberg, Matthew (1987) (ed.) *Readings in Nonmonotonic Reasoning*, Los Altos, CA: Morgan Kaufmann.

Ginsberg, Matthew (1988). Multivalued Logics: A Uniform Approach to Reasoning in AI, to appear in *Computational Intelligence*.

Hanks, Steve and McDermott, Drew (1986). Default Reasoning, Nonmonotonic Logics, and the Frame Problem, *Proc. AAAI-86 1*, pp. 328–333. Reprinted in (Ginsberg 1987).

Hanks, Steve and McDermott, Drew (1987). Nonmonotonic Logic and Temporal Projection, *Artificial Intelligence 33*, pp. 379–412.

Horty, John and Thomason, Richmond (1988). Mixing Strict and Defeasible Inheritance, *Proc. AAAI-88*, pp. 427–432.

Imielinski, Tomasz (1987). Results on Translating Defaults to Circumscription. *Artificial Intelligence 32*, pp. 131–146.

de Kleer, Johan (1986). An Assumption-based TMS. *Artificial Intelligence 28*, pp. 127–162.

Konolige, Kurt (1987). On the Relation between Default and Autoepistemic Logic, in (Ginsberg 1987), pp. 195–226.

Levesque, Hector (1988). All I Know: A Study in Autoepistemic Logic, to appear in *Artificial Intelligence*.

Lifschitz, Vladimir (1985). Computing Circumscription, *Proc. IJCAI-85*, pp. 121–127. Reprinted in (Ginsberg 1987).

Lifschitz, Vladimir (1987). Formal Theories of Action, in Brown, F. (ed.), *The Frame Problem in Artificial Intelligence: Proceedings of the 1987 Workshop*, Morgan Kaufmann, pp. 35–58. Reprinted in (Ginsberg 1987).

Lifschitz, Vladimir (1988). Circumscriptive Theories: A Logic-Based Framework for Knowledge Representation, to appear in the *Journal of Philosophical Logic*.

Lifschitz, Vladimir and Rabinov, Arkady (1988). Miracles in Formal Theories of Action, to appear in *Artificial Intelligence*.

McCarthy, John (1980). Circumscription — a Form of Nonmonotonic Reasoning, *Artificial Intelligence 13*, pp. 27–39. Reprinted in (Ginsberg 1987).

McCarthy, John (1986). Applications of Circumscription to Formalizing Commonsense Knowledge, *Artificial Intelligence 28*, pp. 89–116. Reprinted in (Ginsberg 1987).

Morgenstern, Leora and Stein, Lynn (1988). Why Things go Wrong: a Formal Theory of Causal Reasoning, *Proc. AAAI-88 2*, pp. 518–523.

Moore, Robert (1985). Semantical Considerations on Nonmonotonic Logic, *Artificial Intelligence 25*, pp. 75–94. Reprinted in (Ginsberg 1987).

Morris, Paul (1987). Curing Anomalous Extensions, *Proc. AAAI-87*, pp. 437–442.

Morris, Paul (1988). Autoepistemic Stable Closures and Contradiction Resolution, this volume.

Perlis, Donald (1986). On the Consistency of Commonsense Reasoning, *Computational Intelligence 2*, pp. 180–190. Reprinted in (Ginsberg 1987).

Poole, David (1988). A Logical Framework for Default Reasoning, *Artificial Intelligence 36*, pp. 27–47.

Reiter, Raymond (1978). On Closed World Data Bases, in Gallaire, H. and Minker, J. (eds.), *Logic and Databases*, New York: Plenum Press, pp. 55–76. Reprinted in (Ginsberg 1987).

Reiter, Raymond (1980). A Logic for Default Reasoning, *Artificial Intelligence 13*, pp. 81–132. Reprinted in (Ginsberg 1987).

Reiter, Raymond (1984). Towards a Logical Reconstruction of Relational Database Theory, in M. L. Brodie, J. Mylopoulos and J. W. Schmidt (eds.), *On Conceptual Modelling: Perspectives from Artificial Intelligence, Databases and Programming Languages*, Springer-Verlag, pp. 191–233.

Reiter, Raymond and de Kleer, Johan (1987). Foundations of Assumption-based Truth Maintenance Systems: Preliminary Report, *Proc. AAAI-87*, pp. 183–188.

Ross, Kenneth and Topor, Rodney (1987). *Inferring Negative Information from Disjunctive Databases*, Technical Report 87/1, University of Melbourne; to appear in the *Journal of Automated Reasoning*.