

Stable Closures, Defeasible Logic and Contradiction Tolerant Reasoning*

Paul Morris
IntelliCorp
1975 El Camino Real West
Mountain View, CA 94040

Abstract

A solution to the Yale shooting problem has been previously proposed that uses so-called non-normal defaults. This approach produces a single extension. One disadvantage, however, is that new conflicting information causes the extension to collapse.

In this paper we propose a new formal counterpart to the intuitive notion of a reasonable set of beliefs. The new formalization reduces to the previous one when there are no conflicts. However, when fresh conflicting information is added, instead of collapsing it produces a revised interpretation similar to that obtained by dependency-directed backtracking in a truth maintenance system.

Consideration of the relationship to relevance logic motivates the development of a new formalism for default reasoning, called Defeasible Logic, which behaves like Autoepistemic Logic, but may be more intuitive.

1. Introduction

Recently, much attention has focused on discrepancies between intuition and formalization in nonmonotonic reasoning systems, particularly with respect to the frame problem. For example, in the Hanks-McDermott shooting problem [Hanks and McDermott, 1986], a seemingly natural formalization in terms of Reiter's Default Logic [Reiter, 1980], using normal defaults, supports two interpretations of the events where only one appears to make intuitive sense.

A partial solution to this quandary is proposed in [Morris, 1987], where it is shown that a very similar formulation using a truth maintenance system (TMS) supports only the intuitively sanctioned interpretation.

Moreover, that interpretation is appropriately revised in response to new conflicting information by the mechanism of dependency-directed backtracking.

One drawback of the TMS solution is that truth maintenance is quite limited as an inference mechanism. For example, it is not possible, given justifications $A \rightarrow B$ and $\neg A \rightarrow B$, to conclude B . It is of interest to learn to what extent the inference methods of more powerful logic systems are compatible with intuitively sound nonmonotonic reasoning.

In [Morris, 1987], it is also shown that the non-intuitive interpretation can be excluded within Default Logic by a suitable use of *non-normal* default rules. (In one formulation the frame axiom is replaced by a default rule.) However, in contrast to the TMS behavior, the Default Logic representation does not respond appropriately to new conflicting information. Instead, the addition of such information results in *no* coherent interpretation of the events in the shooting problem.

The difficulty with conflicting new information occurs because non-normal defaults in general are not automatically withdrawn in situations where their application would produce an inconsistency. However, earlier attempts to achieve this behavior by using normal or seminormal defaults have proved untenable because of the countervailing pitfall of unintended interpretation, as illustrated by the shooting problem.

In this paper, we propose a way out of this dilemma by defining a new formal counterpart (within the general framework of autoepistemic logic [Moore, 1985]) to the intuitive notion of an interpretation or reasonable set of beliefs. This new formalization coincides with the previous one in cases where the premises and defaults are free of conflicts. Thus, it can avoid the Hanks-McDermott difficulty by the use of non-normal defaults. Moreover, in a situation where fresh information conflicts with a prior interpretation, the new approach produces one or more revised interpretations which in effect withdraw assumptions as necessary to avoid conflict. The revisions appear to agree well with intuition and to be closely related to those resulting from dependency-directed backtracking in a TMS.

* This research was supported in part by the Defense Advanced Research Projects Agency (DARPA) under contract no. F30602 85 C 0065, and in part by the National Aeronautics and Space Administration and DARPA under contract no. F30602 88 C 0045, through the Rome Air Development Center.

We also consider in this paper the relationship of this new belief revision mechanism to relevance logic (as in [Lin, 1987]). This motivates the development of a new formalism for default reasoning called Defeasible Logic, whose behavior is similar to Autoepistemic Logic, but whose notation and semantics may be closer to our intuition.

2. Background

In Default Logic, an *extension* is the formal counterpart of a reasonable set of beliefs. In the formulation of the shooting example given by Hanks and McDermott there are two extensions, only one of which corresponds to the intended interpretation of the events. We refer to the second extension, corresponding to an unintended (and intuitively unsupported) interpretation, as an *anomalous* extension.

In [Morris, 1987] an example of taxonomic reasoning is presented which is structurally similar to the shooting example, and also produces an anomalous extension. For the purposes of this paper, we will consider a simplified abstract example that also contains the essential structure of the shooting example, but with less extraneous detail. This may be summarized as providing the axioms

1. $\neg abA \supset abB$
2. $\neg abB \supset C$
3. $\neg abA \supset \neg C$

and the normal default inference rules

$$\frac{:\neg abA}{\neg abA} \quad \text{and} \quad \frac{:\neg abB}{\neg abB}$$

As stated, these interact to produce two extensions, just as in the shooting example. In [Morris, 1987], it is shown that a reformulation of the shooting example using a TMS excludes the anomalous extension. In terms of the simplified example, this produces the justifications (the expression $out(X)$ in the left hand side of a justification indicates that X is an OUT justifier)

1. $out(abA) \rightarrow abB$
2. $out(abB) \rightarrow C$
3. $out(abA) \rightarrow \neg C$

which generate a unique well-founded labelling where abB and $\neg C$ are IN. If we subsequently learn that C is true after all, a contradiction occurs which causes dependency-directed backtracking. The contradiction is resolved by introducing a new justification that makes abA come IN, which results in abB and $\neg C$ going OUT.

Thus, the beliefs are appropriately revised in response to new conflicting information.

As in [Morris, 1987], it is possible to use non-normal rules to obtain a Default Logic representation of the simplified example that mirrors the TMS representation in producing a single extension. However, the behavior in response to fresh conflicting information differs from that of the TMS: if C is added as a new axiom, this produces a situation where there is *no* extension.

Konolige [1987] has shown that Moore's Autoepistemic Logic [Moore, 1985] is essentially equivalent in power to Default Logic, with the notion of *stable expansion* playing a role similar to that of an extension. One would expect therefore that a solution to the shooting problem could be obtained using Autoepistemic Logic that would have properties similar to the non-normal default solution for Default Logic. Indeed, Gelfond [1988] has presented an elegant solution that requires only a small alteration to Hanks and McDermott's frame axiom to produce a single stable expansion. However, this shares with the Default Logic solution the difficulty that it collapses in response to fresh conflicting information. Nevertheless, Autoepistemic Logic appears conceptually simpler than Default Logic, and it seems preferable to pursue further developments within this framework.

Following Gelfond, a stable expansion of a set A of axioms may be defined as a solution E to the fixed point equation

$$E = cl[A \cup \{Lx|x \in E\} \cup \{\neg Lx|x \notin E\}]$$

where cl indicates first order logical closure, and we may regard a sentence as an autoepistemic theorem derived from axioms A if it appears in *every* such stable expansion.

In terms of Autoepistemic Logic, the simplified example can be rendered as

$$\begin{aligned} \neg LabA &\supset abB \\ \neg LabB &\supset C \\ \neg LabA &\supset \neg C \end{aligned}$$

With these as axioms there is a unique stable expansion that contains abB but not abA . However, if C is added to the axiom set, there is then no stable expansion because if E were a solution to the fixed point equation we would have

$$\begin{aligned} abA \notin E &\text{ iff } \neg LabA \in E \\ &\text{ iff } \neg C \in E \\ &\text{ iff } abA \in E \end{aligned}$$

which is a contradiction. (Note that from $\neg C$ and C , by classical logic, anything can be deduced, including abA .)

This state of affairs is unsatisfactory because with no

stable expansions, the theory is inconsistent, and the agent can have no formally sanctioned beliefs (or else must believe everything, which is just as bad). This would apply even to sentences that are totally unrelated to the contradiction. This is at odds with respect to our intuitive understanding of commonsense reasoning. This points out an inadequacy in Moore's development of the semantic underpinnings of Autoepistemic Logic. Moore suggests that an ideally rational agent should only consider models in which the agent's beliefs are true. However, if the agent's premises are manifestly inconsistent, it would seem appropriate for the agent to reluctantly concede that some belief(s) might be inaccurate, and instead, perhaps, to assume as many as possible are correct.

A second difficulty is that our intuition from the original shooting example (unfortunately, the abstract example loses some of the force of this) suggests it should not be just the *lack of belief* in abB or abA that implies C or $\neg C$, respectively, but rather the abnormality facts themselves. Intuitively, it seems we should have, for example, $\neg abA \supset \neg C$ (or, equivalently, $C \supset abA$). However, if this is added to the axioms, the anomalous stable expansion reappears.

One possible explanation of this discrepancy (though not a very comforting one) follows from observations of Moore [1985]. These lead to the conclusion that the concept of autoepistemic theorem considered here differs from the usual concept of a theorem in at least one very important respect: in classical logic any theorem can be added to the axioms without altering the theory. This is emphatically not the case for an autoepistemic theorem. For example, given any set of axioms, the sentence $\neg LabA \vee abA$ will be satisfied by every stable expansion. Thus, it is always an autoepistemic theorem. Yet adding it as an axiom in the simplified example will produce a stable expansion containing abA that did not exist previously, causing $\neg LabA$ to lose its status as a theorem. (It is worth noting that the new stable expansion is *strongly grounded*, in the sense of Konolige [1987].) This point suggests the need to exercise caution when axiomatizing a domain: it may be inappropriate to include certain statements as axioms *even though they are intuitively valid*, since their intuitive validity may derive from a status of theoremhood rather than that of axiom.

3. Stable Closure

One possible formal response to the issue of the collapse of the stable expansion is to replace the use of first order logic in the definition of a stable expansion by a relevance or contradiction-tolerant logic (such as in [Lin, 1987]). Thus, cl would indicate closure with respect to a relevance logic. In our example, it would no longer be the case that $\neg C \in E$ iff $abA \in E$, and

there would again be a unique stable expansion, which in this case would contain both C and $\neg C$. The contradiction would still be present, but the damage would be contained so that unrelated beliefs would be unaffected.

This approach is still somewhat unsatisfactory because it does not allow for a revision of belief to resolve the contradiction. Intuitively, it seems reasonable to conclude that abA should be added to the axiom set, whereupon the contradiction disappears. To take another example, a simplification of the well-known "Nixon Paradox" provides the axioms

$$\begin{aligned} \neg LabQ &\supset P \\ \neg LabR &\supset \neg P \end{aligned}$$

which are meant to refer to a specific individual who is both a Quaker and a Republican. Quakers are supposed to be normally pacifists, whereas Republicans are supposed to be typically non-pacifists. Intuitively, the axioms state "If you have no information to suggest the individual is an abnormal Quaker, then conclude he is a pacifist" and "If you have no reason to think the individual is an abnormal Republican, then conclude he is a non-pacifist." With cl as first order closure, there is no stable expansion, so that our previous remarks hold about the unfortunate effects for additional unrelated beliefs. With relevance closure, a stable expansion that contains both P and $\neg P$ is obtained, so that the damage is mitigated. However, relevance closure still does not allow the inference $abQ \vee abR$ (i.e., "Nixon is either an abnormal Quaker or abnormal Republican"), which seems intuitively valid.

From a semantic point of view, use of relevance logic amounts to an acknowledgement that some of one's beliefs may be in error. A seemingly less radical position would be to reluctantly accept, in the face of convincing evidence, that one's base beliefs may be *incomplete*. This approach leads to a somewhat different formalization that, as we will see, meets the difficulty discussed above. (Later on, we will see that the formalism is related to a form of relevance logic, after all.)

Proceeding along these lines, we first define an *expandable* set of axioms to be one that possesses a stable expansion. An *augmentation* of a set of axioms A is a set B such that $B = cl[A \cup G]$ where cl is first order logical closure, and G is a set (possibly empty) of ordinary sentences of the language (i.e., not involving the L operator). Next we define a *stable completion* of a set of axioms A to be a minimal expandable augmentation of A . Finally, we define a *stable closure* of a set A to be a stable expansion of a stable completion of A .

Thus, a stable closure of an axiom set A is a solution E of the fixed point equation

$$E = cl[B \cup \{Lx|x \in E\} \cup \{\neg Lx|x \notin E\}]$$

where B is a *minimal augmentation* of A for which the equation has a solution. (The minimality condition means that if B' is another augmentation of A for which the equation has a solution, and if $B' \subseteq B$, then $B' = B$.) We say the stable closure is *generated* by any set G of ordinary sentences such that $cl[A \cup G] = B$.

(The reason why the definition of augmentations calls for them to be logically closed is that otherwise we could “cheat” on the minimality condition by using a conjunct $P \wedge Q$ in B where P alone would have sufficed.¹)

We now replace stable expansion by stable closure as the formal counterpart of a reasonable set of beliefs based on a set of axioms, and we redefine an *autoepistemic theorem* to be a sentence which is contained in all stable *closures* of the axioms.

It is an immediate consequence of the definition that the concept of stable closure reduces to that of stable expansion when the latter exists. (Note that for any axiom set T , the stable expansions of $cl[T]$ are the same as those of T .) For our simplified version of the shooting example

$$\begin{aligned} \neg LabA &\supset abB \\ \neg LabB &\supset C \\ \neg LabA &\supset \neg C \end{aligned}$$

there is therefore a unique stable closure that coincides with the stable expansion considered earlier. Note that if we now add C as a new axiom, there is no longer a stable expansion, but there *is* a stable closure, generated by $\{abA\}$. (To see this, observe that after adding C , any augmentation of the axioms that does *not* contain abA will be non-expandable, since the conflict will still be present. However, the addition of abA is sufficient to produce an expandable set.) This stable closure contains $\neg LabB$ and C instead of $\neg LabA$ and $\neg C$, i.e., it produces the revisions we expect intuitively.

It is important to note that the concept of stable closure of A differs from the concept of minimal stable set containing A . For instance, in the above example before adding C , there is a unique stable closure. However, it is easy to verify that there are two minimal stable sets that contain the axioms. (One is the usual stable expansion; call it E . To see there is another, let S be a stable expansion of the set obtained by adding abA to the axioms. Then S is also a stable superset of the original axioms. Note that it is not a superset of E .) The stable closure also differs from the concept of minimal stable expansion, as considered in [Konolige,

1987], since, as we have seen, a stable closure may exist even when there is *no* stable expansion.

Now consider once again the simplified Nixon example

$$\begin{aligned} \neg LabQ &\supset P \\ \neg LabR &\supset \neg P \end{aligned}$$

Here there are two stable closures, generated by $\{abQ\}$ and $\{abR\}$, respectively. Thus, the sentence $abQ \vee abR$ will be satisfied by every stable closure, and so, it is an autoepistemic theorem.

These revisions appear to agree well with those obtained from dependency-directed backtracking. (To compare Autoepistemic Logic to a TMS, we can draw a rough analogy between $\neg LX$ and $out(X)$.) In particular, with LA as a solitary axiom, which is the autoepistemic analogue of $out(A) \rightarrow false$, we get a stable closure containing A , whereas with $LA \vee LB$, which is the analogue of $out(A) \wedge out(B) \rightarrow false$, we get two stable closures with A and B , respectively.

One area where the stable closure gives a result unexpected from the TMS point of view is with the axiom $LA \vee A$. This is the analogue of the TMS expression $out(A) \rightarrow A$, which is the archetypal odd loop. However, a nontrivial stable closure does exist for this, containing A . Thus, $out(A) \rightarrow A$ behaves similarly to $out(A) \rightarrow false$.

Note that irrespective of the conflicts that exist among the defaults, there is always at least one expandable augmentation of the axioms, namely, the set of all sentences. There are cases in which this is the (sole) stable completion, for example, with the axiom set $\{La, L\neg a\}$. It is an open question whether every axiom set possesses a stable completion. (It is conceivable that the class of expandable augmentations of some set has infinite descending chains, with no minimal elements.)

4. Defeasible Logic

We wish to explore the relationship between stable closures and relevance logic. This motivates us to look for an approach in which ordinary axioms can be represented as defaults, so that the resolution of conflicts provided by the stable closure supports a form of contradiction tolerant reasoning.

In furtherance of this approach, we first consider a syntactic variant of Autoepistemic Logic in which defaults are more conspicuous. Observe that sentences of the form $\neg LX$ provide the basic source of defaults in Autoepistemic Logic. We make this more apparent by defining

$$D = \{\neg LX | X \in U\}$$

where U is the set of all sentences in the language. We define a mapping $R : D \rightarrow U$ by

¹ I am grateful to Matt Ginsberg for pointing out this defect in an earlier definition.

$$R(\neg LX) = X$$

for any $\neg LX \in D$. We will abbreviate $R(d)$ by the notation $\bullet d$ and verbalize it as "Revoke d ." We call " \bullet " the *revoke* operator.

We can now express a stable expansion as a solution of the fixed point equation

$$E = cl[A \cup \{d | d \in D \text{ and } \bullet d \notin E\} \\ \cup \{\neg d | d \in D \text{ and } \bullet d \in E\}]$$

So far this is merely syntactic sugar. However, the new notation suggests a quite different semantics where the defaults are regarded as objective rather than subjective facts. For instance, the simplified Nixon example can be represented as

$$dQ \supset P \\ dR \supset \neg P$$

where dQ and dR belong to D . Instead of thinking of dQ as a subjective statement about one's beliefs, it is tempting to interpret it as representing the objective statement "Nixon is a normal Quaker," which is true by default.

With this perspective we can abandon the autoepistemic origin of the default set D and allow it to be any set of ordinary sentences. In this approach, "revoke" is a modal operator whose meaning derives solely from its use in the fixed point equation. This gives us a new formalism for default reasoning which is somewhat different from Autoepistemic Logic. We name it *Defeasible Logic*.

In Defeasible Logic, the defaults may be viewed as a collection of additional "axioms" which are tentative in the sense that they are subject to being revoked. These "axioms" may take the form of the negations of abnormality propositions. Thus, our simplified shooting example might be represented as

$$\neg abA \supset \bullet \neg abB \\ \neg abB \supset C \\ \neg abA \supset \neg C$$

where $\neg abA$ and $\neg abB$ are members of the default set. Note the use of the revoke operator to exclude the unintended interpretation.

It is interesting to recall our earlier remarks concerning the intuitive validity of $\neg abA \supset \neg C$ and to observe that in this approach it appears as a full-fledged axiom. Note, however, that in this framework, $\neg abA \vee \bullet \neg abA$ holds in every stable expansion, but may not be added as an axiom without altering the theory.

Rather than use abnormality propositions, it is possible to introduce implications directly as defaults. For example, the sentence $Bird \supset Fly$ could be a default, where $Bird$ denotes that a particular individual, say, Tweety, is a bird, and Fly that Tweety

can fly. We can then allow for the possibility of Tweety being exceptional by virtue of being an ostrich, by adding the axiom

$$Ostrich \supset \bullet (Bird \supset Fly)$$

where $Ostrich$ denotes Tweety is an ostrich. Note, however, that with this representation, a stable expansion that contains $Ostrich$ will also contain $\neg (Bird \supset Fly)$, so that $\neg Fly$ can be deduced. This deduction could be avoided by an abnormality approach.

There is a sense in which Defeasible Logic and Autoepistemic Logic are equivalent in power. A Defeasible Logic theory may be simulated by an Autoepistemic Theory by introducing an axiom $\neg L\bullet d \equiv d$ corresponding to each default d . Here, $\bullet d$ is regarded as an atomic proposition. Conversely, an Autoepistemic Theory may be simulated by a Defeasible Logic Theory by introducing a default $\neg LA$, and an axiom $A \supset \bullet \neg LA$, corresponding to each sentence A . In this case, LA is regarded as an atomic proposition.

Some variations on the formalism of Defeasible Logic may be worth exploring. For example, it seems intuitively desirable to have $\neg \bullet d$ hold in situations where d has not been revoked. We could achieve this by replacing the original fixed point equation with

$$E = cl[A \cup \{d \wedge \neg \bullet d | d \in D \text{ and } \bullet d \notin E\} \\ \cup \{\neg d | d \in D \text{ and } \bullet d \in E\}]$$

which is easily seen to be equivalent to

$$E = cl[A \cup \{\bullet d \supset \neg d | d \in D\} \\ \cup \{d | d \in D \text{ and } \bullet d \notin E\} \\ \cup \{\neg d | d \in D \text{ and } \bullet d \in E\}]$$

We can effect a further simplification by restricting our attention to solutions that are *strongly grounded* in a sense analogous to that of Konolige [1987]. We define a solution E of the above equation to be strongly grounded if it satisfies

$$E \subseteq cl[A \cup \{\bullet d \supset \neg d | d \in D\} \\ \cup \{d | d \in D \text{ and } \bullet d \notin E\}]$$

For strongly grounded solutions, the fixed point equation further reduces to

$$E = cl[A \cup \{\bullet d \supset \neg d | d \in D\} \\ \cup \{d | d \in D \text{ and } \bullet d \notin E\}]$$

Note that all the solutions of this last equation are strongly grounded.

4.1. Stable Closure and Relevance Logic

The concept of stable closure can also be applied to Defeasible Logic. In this case, an *augmentation* of an axiom set A is a set $cl[A \cup G]$ where G has the form $\{\bullet d | d \in D1\}$ for some subset $D1$ of the default set D .

The definitions of *stable completion*, *stable closure*, and *generator* of a stable closure are as before. We now compare the resolution of inconsistencies provided by the stable closure with that of a relevance logic.

Consider an ordinary first-order theory. This is equivalent to a Defeasible Logic theory with the same axioms, and with an empty set of defaults. Since the axioms come from the first-order theory, none of them will involve the revoke operator. Notice that if the axioms are consistent, replacing each axiom by an equivalent default leaves the essential theory unchanged, with each ordinary theorem becoming a Defeasible Logic theorem.

It is, however, when the axioms are inconsistent that the transformation produces an interesting result: by considering stable closures generated by sentences of the form $\bullet A$, where A is among the original axioms, we arrive at a nontrivial definition of the concept of "theorem" within an inconsistent system.

For example, with the axiom set $\{A, B, \neg AV \neg B\}$, the stable closures are generated by $\{\bullet A\}$, $\{\bullet B\}$, and $\{\bullet(\neg AV \neg B)\}$, respectively. Thus, the stable closures consist of the sets $cl\{B, \neg AV \neg B\}$, $cl\{\neg AV \neg B, A\}$, and $cl\{A, B\}$, respectively, and the theorems lie in the intersection of these three sets.

This approach simulates a relevance logic in the sense of Lin [1987]. Actually, it's not hard to see that it coincides with one of the formulations discussed by Lin: a sentence is a theorem in this approach if it is entailed by every maximal consistent subset of the axioms. Lin cites as a disadvantage of this definition the fact that the resulting set of theorems is not recursively enumerable; unfortunately, this is part of the price that must be paid for a nonmonotonic reasoning system.

5. Conclusions

We have presented an approach to formalizing the intuitive notion of a reasonable set of beliefs that on the one hand allows us to express a preference for one interpretation over another, thus solving the shooting problem, and on the other hand allows us to revise an interpretation in a reasonable way to account for new conflicting information.

At first sight, it may appear that the concept of stable closure is likely to have unappealing computational properties. However, the computational model we have in mind is that of truth maintenance which appears to compete favorably with other approaches. The intention is to provide a better formal basis for exploring TMS-like mechanisms. For example, the notion of stable closure appears to capture at least part of the idea of dependency-directed backtracking.

We have also examined the relationship of the approach to the kind of contradiction-tolerant reasoning considered in relevance logic. This motivated

the creation of a new formalism for default reasoning, called Defeasible Logic, which has behavior similar to that of Autoepistemic Logic, but which may match our intuition better in some respects.

Many issues remain. There may be a lack of convincing linguistic evidence for the reality of a revoke operator in commonsense reasoning, compared to, say, that for a belief operator. We mentioned earlier that one odd feature of Autoepistemic Logic is the fact that theorems can not in general be added to the axioms without the possibility of altering the theory. The same is true of Defeasible Logic. It would be nice if a formalism could be found that avoided this, while preserving the attractive properties of the existing formalisms. Finally, further investigations are needed to develop the relationship between truth maintenance and formal systems of default reasoning more precisely.

Acknowledgements

I wish to thank Bob Nado and Matt Ginsberg for useful discussions on the subject of this paper. I also thank the referees for their criticisms, which hopefully have led to improvements in this paper.

References

- [Gelfond, 1988] Gelfond, M. Autoepistemic Logic and Formalization of Commonsense Knowledge. Preprint, Computer Science Dept., Univ. of Texas at El Paso. 1988
- [Hanks and McDermott, 1986] Hanks, S., and McDermott, D. Default reasoning, nonmonotonic logics, and the frame problem. In *Proceedings AAAI-86*, pages 328-333. Philadelphia, 1986.
- [Konolige, 1987] Konolige, Kurt. On the Relation between Default Theories and Autoepistemic Logic. In *Proc. IJCAI-87*. Milan, Italy, 1987.
- [Lin, 1987] Lin, Fangzhen. Reasoning in the presence of inconsistency. In *Proceedings AAAI-87*, pages 139-143. Seattle, 1987.
- [Moore, 1985] Moore, R.C. Semantical Considerations on Nonmonotonic Logic. *Artificial Intelligence* 25:75-94, 1985.
- [Morris, 1987] Morris, P.H. Curing Anomalous Extensions. In *Proceedings AAAI-87*, pages 437-442. Seattle, July, 1987.
- [Reiter, 1980] Reiter, Raymond. A logic for default reasoning. *Artificial Intelligence* 13:81-132, 1980.