Curing Anomalous Extensions

Paul Morris

IntelliCorp
1975 El Camino Real West
Mountain View, CA 94040

Abstract

In a recent paper, Hanks and McDermott presented a simple problem in temporal reasoning which showed that a seemingly natural representation of a frame axiom in nonmonotonic logic can give rise to an *anomalous* extension, i.e., one which is counter-intuitive in that it does not appear to be supported by the known facts.

An alternative, less formal approach to nonmonotonic reasoning uses the mechanism of a *truth maintenance system* (TMS). Surprisingly, when reformulated in terms of a TMS, the anomalous extension noted by Hanks and McDermott disappears. We analyze the reasons for this. First it is seen that anomalous extensions are not limited to temporal reasoning, but can occur in simple non-temporal default reasoning as well. In these cases also, the natural TMS representation avoids the problem. Exploring further, it is observed that the form of the TMS justifications resembles that of *nonnormal* default rules. Nonnormal rules have already been proposed as a means of avoiding anomalous extensions in some non-temporal reasoning situations. It appears that, suitably formulated, they can exclude the anomalous extension in the Hanks-McDermott case also, although the representation does not adjust smoothly to fresh information, as does the TMS. Some variant of nonnormal default appears to be required to provide a correct semantic basis for truth maintenance systems.[1]

I. Introduction

One of the central requirements for an effective temporal reasoning system is a reasonable solution of the *frame problem*. The frame problem is that of specifying the effects of actions in a way that allows efficient determination of the properties that hold in subsequent states. In particular, a representation is needed that allows exploitation of the fact that in many situations of interest most properties are unchanged by a given action.

The development of nonmonotonic and default logics has been seen as promising a way of achieving such a representation within a well-understood declarative framework. Unfortunately, a foundational difficulty in this approach has recently been uncovered by Hanks and McDermott [Hanks and McDermott, 1986], who present an example of temporal reasoning where the natural default logic representation is shown to be inadequate for deriving some intuitively sound conclusions.

Shoham [Shoham, 1986] proposes a solution to this difficulty which uses a default reasoning mechanism specific to temporal reasoning. This seems to suggest that the problem is an artifact of temporal reasoning. To counter this view, and show that the problem is a wider one for nonmonotonic logic, we will present an example drawn from non-temporal default reasoning that reproduces the difficulty.

An alternative approach to default reasoning that has seen considerable use in practice, but has undergone relatively little formal study involves the mechanism of a *truth maintenance system* (TMS). We will see that, surprisingly, the difficulty noted for nonmonotonic logic disappears when the example is reformulated in terms of a truth maintenance system. Moreover, the TMS revises its beliefs appropriately in response to fresh information.

An examination of the TMS representation suggests the use of nonnormal default rules in the Hanks-McDermott example. Nonnormal rules, suitably formulated, can indeed exclude the anomalous extension. However, in contrast to the TMS, the nonnormal rule representation does not respond smoothly to changes in belief. The difference arises because in cases where the TMS produces a contradiction, provoking backtracking, the default rule representation can result in *no* extensions. A small change is suggested to the semantics of default rules to make them more closely approximate the behavior of TMS justifications.

With the suggested modification, the use of nonnormal default rules opens up the possibility of having inconsistent extensions, even though the underlying monotonic theory is consistent. Rather than place the burden of excluding inconsistencies on the default logic mechanism, one might regard applications of *reductio ad absurdum* reasoning to resolve inconsistencies as an extra-logical operation that modifies the existing axiom set. In support of this view, we present an example from the area of planning which suggests that such reasoning needs to make distinctions at a level beyond ordinary logic.

II. Hanks-McDermott Anomaly

Common approaches to nonmonotonic logic use default inference rules [Reiter, 1980] or circumscription [McCarthy, 1980] to extend a set of beliefs with as many default assumptions (and their deductive consequences) as can be consistently added. The resulting larger set of beliefs is called an *extension*. Note, however, that the relative consistency of defaults may depend upon the order in which they are added, giving rise to multiple competing extensions. One resolution of this (suggested by Hanks and McDermott) is to regard a statement as being a nonmonotonic "theorem" if it holds in *every* extension.

[1] This research was supported in part by the Defense Advanced Research Projects Agency under contract No. F30602 85 C 0065. The views expressed are those of the author and do not necessarily represent the position or policy of DARPA.

Hanks and McDermott present an example of temporal reasoning (the "shooting example") that gives rise to multiple extensions. However, the example is such that only one of these corresponds to our intuition. A second extension corresponds to a possible interpretation of the events, but one that intuitively is not supported by the known facts. This means that some conclusions which are intuively valid can not be derived as theorems.

In the shooting example, a sharpshooter loads his gun and lies in wait for a victim. When the victim appears, the gunman shoots. Assuming the gun is loaded at the time the shot occurs, the victim dies. Thus, we have a state $S1$ at which the gun is loaded, followed by a waiting period until another state $S2$ where the victim is alive, followed by a shooting action resulting in a third state $S3$. The question is whether the victim is alive or dead at $S3$. In the anomalous extension the gun mysteriously becomes unloaded during the waiting period, so that the victim survives.

Following Hanks and McDermott, we use the notation $T(f,s)$ to indicate fact f is true in state s. For each such proposition and action e, we have a *frame axiom* of the form

$$T(f, s) \land \neg AB(f, e, s) \supset T(f, RESULT(e,s))$$

where $RESULT(e,s)$ represents the state resulting from applying action e to state s, and $\neg AB(f, e, s)$ represents the assertion that f is unaffected by e in state s.

We can summarize the shooting example as providing the axioms

1. $T(ALIVE, S2)$
2. $T(ALIVE, S2) \land \neg AB(ALIVE, SHOOT, S2) \supset T(ALIVE, S3)$
3. $T(LOADED, S2) \supset AB(ALIVE, SHOOT, S2)$
4. $T(LOADED, S2) \supset T(DEAD, S3)$
5. $T(LOADED, S1)$
6. $T(LOADED, S1) \land \neg AB(LOADED, WAIT, S1) \supset T(LOADED, S2)$

and the default rules

$$\frac{:M\, \neg AB(LOADED, WAIT, S1)}{\neg AB(LOADED, WAIT, S1)}$$

$$\frac{:M\, \neg AB(ALIVE, SHOOT, S2)}{\neg AB(ALIVE, SHOOT, S2)}$$

As Hanks and McDermott point out, with the above axioms, each default rule defeats the other. This means that there is an extension where the second default is applied, but not the first, leading to the conclusion $T(ALIVE, S3)$, which intuitively does not appear to be supported by the facts presented.

It is not obvious from this example where the difficulty lies. At first sight it appears conceivable that temporal reasoning possesses special problems related to the flow of time and causality. Thus, Shoham proposes that chronologically earlier defaults should be added first when constructing an extension. We argue here that to pin the blame on temporal reasoning is to misdiagnose the disease; the real difficulty is independent of temporal reasoning. Indeed, other authors [Reiter and Criscuolo, 1983, McCarthy, 1986] (see also [Etherington, 1987]) have previously observed anomalous extensions arising in non-temporal reasoning. To further clarify the situation, we present a natural example from non-temporal default reasoning that reproduces the difficulty and is structurally very similar to the shooting example.

Consider the following statements.

- Animals normally can not fly.
- Winged animals are exceptions to this.
- All birds are animals.
- Birds normally have wings.

Now suppose Tweety is a bird. Given the above statements, one would intuitively conclude that Tweety has wings. One would NOT conclude that Tweety is unable to fly, although Tweety is obviously also an animal. However, just as in the Hanks-McDermott shooting example, a seemingly natural representation of this example in nonmonotonic logic produces two extensions. In one of these, Tweety is a normal bird that has wings and may be able to fly. This matches our intuition. In the other extension, Tweety is a normal animal, but an abnormal bird. Thus, he is wingless, and unable to fly. This is an anomalous extension.

To formalize the example, we make the assignments

```
A = "Tweety is an animal"
B = "Tweety is a bird"
W = "Tweety has wings"
G = "Tweety can not fly (is Grounded)"

abA = "Tweety is an abnormal animal with respect
       to not flying"
abB = "Tweety is an abnormal bird with respect
       to having wings"
```

The negations of the last two statements correspond to defaults.

Continuing the formalization, we have the implications

1. $A \land \neg abA \supset G$
2. $W \supset abA$
3. $B \supset A$
4. $B \land \neg abB \supset W$

We are also assuming B as an axiom. Following the analogy with the shooting example, we have the default inference rules

$$\frac{:M\, \neg abA}{\neg abA}$$

and

$$\frac{:M\, \neg abB}{\neg abB}$$

The first of these can be applied to construct an extension that includes ¬abA. However, we can then use implication 2 to deduce ¬W. This allows us to conclude abB by 4, which prevents the second default rule from being applied. Thus, the extension does not include ¬abB. This extension is counter-intuitive. It is easy to verify that the intuitive extension can also be obtained.

This shares with the shooting example the characteristic that by assuming a normality, we could deduce an abnormality, and thereby arrive at an anomalous extension. In order to isolate the problem further, we will try to remove extraneous features from the examples, and boil them down to their essential elements. In the shooting example, we could simplify the implications by omitting undisputed facts such as $T(LOADED, S1)$. This gives us

$$\neg AB(LOADED, WAIT, S1) \supset T(LOADED, S2)$$

$$T(LOADED, S2) \supset AB(ALIVE, SHOOT, S2)$$

$$T(LOADED, S2) \supset T(DEAD, S3)$$

$$\neg AB(ALIVE, SHOOT, S2) \supset T(ALIVE, S3)$$

The intermediate fact $T(LOADED, S2)$ is significant only in connecting the two abnormality facts. Eliminating this "middleman" gives a further simplification to

$$\neg AB(LOADED, WAIT, S1) \supset AB(ALIVE, SHOOT, S2)$$

$$\neg AB(LOADED, WAIT, S1) \supset T(DEAD, S3)$$

$$\neg AB(ALIVE, SHOOT, S2) \supset T(ALIVE, S3)$$

Applying a similar process of simplification to the bird example gives

$$\neg abB \supset abA$$

$$\neg abA \supset G$$

The pattern in both cases is that we have two defaults A and B such that $A \supset \neg B$. Logically, this is equivalent to $B \supset \neg A$, so that the crucial information would appear to be symmetric in A and B. Nevertheless, in the examples seen, the intuitively correct extension includes A but not B.

It might be noted that the anomalous extension in the bird example could be ruled out by including $B \supset abA$ as an additional axiom. However, this fails to accurately capture the knowledge that it is wings that make flight possible for birds. In particular, if we subsequently learned that Tweety had his wings torn off in an accident, we would have no way to revert to the default for animals and conclude that Tweety is unable to fly.

III. Truth Maintenance

Truth maintenance systems were introduced by Doyle [Doyle, 1979] and have been refined and extended by many workers since then. The TMS model we use in this paper is essentially that of Doyle.

Truth maintenance systems have often been regarded as performing a kind of resource bounded inference rather than true inference in that beliefs are propagated based on the current status of other beliefs rather than their ultimate status. As a practical matter, however, a TMS is usually left alone until it reaches a quiescent state. Such a state corresponds to a fixed point, just as in default logic.

In a truth maintenance system, the notion of default does not arise directly. Instead, one is allowed to use nonmonotonic *justifications*. These may be regarded as inferences which may partly depend upon a state of ignorance with respect to certain facts. For example, if one's car is left parked, and there is no reason to think it is moved, one expects it to be there when one returns. We will write out(A) to represent ignorance of fact A. Then the car example could be expressed by the justification

$$P \wedge out(M) \rightarrow R$$

where P denotes the car is parked, M that it is moved, and R that it will still be there upon return. We call P an *IN justifier* of the justification. We say M is an *OUT justifier*.[2] One may view an OUT justifier as providing a kind of built-in default.

[2] This notation for nonmonotonic justifications differs somewhat from the standard one which segregates the IN and OUT justifiers, and writes them separately.

In terms of a TMS, the most natural coding of the bird example would be

1. $A \wedge out(abA) \rightarrow G$
2. $W \rightarrow abA$
3. $B \rightarrow A$
4. $B \wedge out(abB) \rightarrow W$

An unconditional premise, such as B in this example, is represented by a justification with an empty set of justifiers.

In a TMS the notion of a *labelling* plays a role similar to that of an extension in nonmonotonic logic. One might define a valid labelling as an assignment of IN/OUT status to each of the propositions in such a way that each of the justifications is satisfied. For example, if the labelling specified A as IN and abA as OUT, then to satisfy the first justification, G would have to be IN. If the labelling specified B as IN and W as OUT, then to satisfy the fourth justification, abB would have to be IN. With this definition there would be two valid labellings corresponding to the two extensions discussed earlier. In particular, there would be an anomalous labelling having abB labelled as IN.

However, the label propagation mechanisms employed in truth maintenance systems generally have the property that only *well-founded* labellings are obtained. A labelling is well-founded if it is valid in the sense above and, in addition, every proposition labelled IN is *well-justified*, that is, it is the conclusion of some justification whose OUT justifiers are all labelled OUT, and whose IN justifiers are all themselves well-justified. Since abB is not the conclusion of any justification, there is no *well-founded* labelling that labels it as IN. Thus, for a truth maintenance system, there is a single accessible extension in this example, namely the one corresponding to our intuition.

It is interesting to note that if we subsequently learn that W is false, the contradiction-handling machinery of a truth maintenance system (in what might be considered an application of *reductio ad absurdum*) will install a new "backward" justification for abB. This causes the second extension to become accessible since it now corresponds to a well-founded labelling. This is exactly what we expect intuitively: if we learn that due to an unfortunate accident poor Tweety is wingless, we indeed want to revert to the default for animals, and conclude he is unable to fly.

The natural TMS representation of the implications in the shooting example is

$$T(ALIVE, S2) \wedge out(AB(ALIVE, SHOOT, S2)) \rightarrow T(ALIVE, S3)$$

$$T(LOADED, S2) \rightarrow AB(ALIVE, SHOOT, S2)$$

$$T(LOADED, S2) \rightarrow T(DEAD, S3)$$

$$T(LOADED, S1) \wedge out(AB(LOADED, WAIT, S1)) \rightarrow T(LOADED, S2)$$

which similarly excludes the anomalous labelling since $AB(LOADED, WAIT, S1)$ is not the conclusion of any justification. Notice that if we subsequently learn $T(ALIVE, S3)$, a contradiction is produced, causing backtracking. The only possible culprit is $out(AB(LOADED, WAIT, S1))$, so the TMS installs a justification for $AB(LOADED, WAIT, S1)$, which causes a shift to the second extension. Again, this satisfies our intuitive expectations: if we learn the victim survives, the only possibility given the statement of the problem is that the gun became unloaded.

In both examples, applying the process of simplification considered earlier produces a pattern of the form out(A) $\rightarrow$ B and out(B) $\rightarrow$ C. In general, with a pattern of this form where there are no cycles, there is a single well-founded labelling. Moreover, a TMS will arrive at that labelling irrespective of the order in which the justifications are added.

IV. Nonnormal Defaults

The question arises: what property of a TMS enables it to escape the Hanks-McDermott anomaly? It might appear at first sight that default logic is unable to capture the well-foundedness requirement, or that the limited nature of the inference performed by the TMS is responsible. The first possibility can be ruled out because the minimality requirement for extensions in default logic is there to ensure well-foundedness. The unidirectional nature of TMS inference does play a role. However, we will see that it is possible to exclude the anomaly by changing the default logic representation in a way suggested by the TMS formulation.

If we examine the behavior of an OUT justifier B in a TMS justification

$$A \wedge \text{out}(B) \rightarrow C$$

we see that it is satisfied when B is not IN. For a quiescent state of the TMS, this means B is not derivable. In a logic system the non-derivability of B would be equivalent to $\neg$B being consistent with the other facts. This suggests that we regard the entire TMS justification as a default inference rule of the form

$$\frac{A : \mathbf{M}\,\neg B}{C}$$

Taking this approach in the bird example, we replace justifications 1 and 4 by the default rules

$$\frac{A : \mathbf{M}\,\neg \text{abA}}{G}$$

and

$$\frac{B : \mathbf{M}\,\neg \text{abB}}{W}$$

respectively. Now observe that application of the second default rule (together with justification 2) defeats the first default rule, but not vice versa, so we end up with a single extension.

In the shooting example, the frame axioms, instead of being implications, become default rules of the general form

$$\frac{T(f,s) : \mathbf{M}\,\neg AB(f,e,s)}{T(f,RESULT(e,s))}$$

It may be verified that this formulation eliminates the anomalous extension.

We see in both examples the use of so-called *nonnormal* default rules. Hanks and McDermott did not discuss nonnormal defaults in their paper.

Reiter and Criscuolo [Reiter and Criscuolo, 1983] (also [Etherington, 1987]) suggest the use of a special kind of nonnormal default, called a *seminormal* default, to exclude anomalous extensions. These have the general form

$$\frac{\alpha : \mathbf{M}\,(\beta \wedge \gamma)}{\beta}$$

and are obtained by modifying normal defaults to anticipate conditions which would render them inappropriate. Thus, in the bird example one might use the seminormal rules

$$\frac{A : \mathbf{M}\,(G \wedge \neg \text{abA})}{G}$$

and

$$\frac{B : \mathbf{M}\,(W \wedge \neg \text{abB})}{W}$$

in place of the nonnormal rules suggested earlier. This approach excludes the anomalous extension in the bird problem as it stands. Unfortunately, if we make the bird example more closely resemble the shooting example by including

$$W \supset \neg G$$

as an additional axiom, the anomalous extension is restored (this would not be the case for the original nonnormal rules). To see this, note that application of the first rule, in combination with the new axiom, gives $\neg$W, which defeats the second rule. It may be verified that the corresponding formulation of the shooting problem using seminormal rules also produces the anomalous extension. Thus, even *semi*normal rules appear insufficient to properly constrain interactions between defaults in all cases.

Although the more general nonnormal default rules that we have considered *do* exclude the anomalous extension, as they stand they do not respond appropriately to the addition of new information (in contrast to the TMS). Consider in the bird example what happens if we learn that Tweety is missing his wings, so that $\neg$W becomes a new axiom. Rather than produce an inconsistent extension, as would the TMS, the nonnormal rules interact in such a way as to produce *no* extension. To see this, observe that with $\neg$W as an axiom, the default rule

$$\frac{B : \mathbf{M}\,\neg \text{abB}}{W}$$

defeats *itself*, since from W and $\neg$W, using ordinary (classical) deduction, one can derive abB. Indeed, it has been proved [Reiter, 1980] that a default logic extension is inconsistent if and only if the underlying monotonic theory is itself inconsistent.

In view of the difficulty in responding to new information, it might be preferable if Default Logic somehow emulated a TMS and allowed the possibility of inconsistent extensions. We make an informal suggestion here as to how this might be achieved. Suppose the interpretation of **M** B in the default rule

$$\frac{A : \mathbf{M}\,B}{C}$$

is changed from the usual "it is consistent to assume B" to "it is consistent to assume B *or B is provable.*" Now an application of a default rule that results in an inconsistency will not automatically undercut itself, so the possibility arises of having an inconsistent fixed point.

V. Isolated Defaults

It seems a little drastic to represent all justifications as default rules. Such rules are unidirectional and we would like to preserve as much as possible of the bidirectional nature of implicational inference (e.g., if $\neg A \supset B$ then $\neg B \supset A$). Intuitively, the entire default content of a justification resides in the OUT justifiers. In the car example considered earlier, there appears to be an underlying assumption that the car will not be moved which is separable from the other parts of the justification. Indeed, one of the attractive aspects of normal defaults was that they could be isolated in this way. We now consider whether a similar isolation can be achieved for nonnormal defaults.

Consider the expression "out(X)." One way of viewing this is that it represents a proposition in its own right distinct from X, although related to it. To achieve behavior resembling that of a truth maintenance system, we could add

$$\frac{:M\, \neg X}{\text{out}(X)}$$

as a default logic rule. With this separate specification of the OUT justifiers, we can represent justifications as ordinary implications. Looking once again at the bird example, we can show that with this formulation there is no extension containing out(abA). Suppose that there is such an extension. Since it is NOT the case that out(abA) $\rightarrow$ ¬abA, there is no direct conflict between out(abA) and ¬abB. Thus, the default rule for out(abB) can be applied to conclude out(abB). But this allows us to deduce abA, which prevents the default rule for out(abA) from being used. Thus, the extension does not contain out(abA) after all. Observe that there is no difficulty with the extension containing out(abB).

One cautionary note is in order with respect to the suggested default logic formalization of out(X). It has been customary in truth maintenance systems to represent the adoption of A as an assumption by introducing

$$\text{out}(\neg A) \rightarrow A$$

as a justification. If we have a similar justification

$$\text{out}(\neg B) \rightarrow B$$

for B, and also have $A \rightarrow \neg B$, then the default logic formalization once again has multiple extensions. In this case the TMS has only a single well-founded labelling because, from the TMS point of view, the inference $A \rightarrow \neg B$ is unidirectional. Assuming the single extension is what is intended, it appears unfortunate to rely on unidirectionality with respect to a *monotonic* justification like $A \rightarrow \neg B$ to achieve it in the TMS. This would mean that monotonic justifications also would have to be represented as inference rules, rather than implications, in the default logic formalization. To avoid this difficulty, we suggest using justifications

$$\text{out}(\sim A) \rightarrow A$$
$$\text{out}(\sim B) \rightarrow B$$
$$A \rightarrow \sim B$$

to represent the intended situation, where $\sim X$ is a dummy proposition distinct from X or $\neg X$ (one might read $\sim X$ as "X is defeated"). Now there is a single extension in the suggested default logic formalization.

With the $\text{out}(\sim D) \rightarrow D$ representation of defaults, it is possible to derive inconsistencies. For example, if we have defaults A and B defined analogously to D, and if $A \rightarrow \neg B$, then both justifications are operative, so that a contradiction is derived. We argue that inconsistency in this kind of situation is preferable to having multiple extensions. Since monotonic inferences can give rise to inconsistencies, it seems reasonable to allow nonmonotonic inferences the same privilege. The inconsistency can be removed by supplying a new justification such as $A \rightarrow \sim B$ or $B \rightarrow \sim A$. We argue that the choice between these, or other resolutions, is best left to an extra-logical contradiction-handling procedure. Indeed, we will see that in any case contradiction-handling needs to be sensitive to extra-logical issues.

It is worth remarking that the well-known result [Charniak *et al.*, 1979] for truth maintenance systems, that the absence of odd loops guarantees the existence of a well-founded labelling, appears related to the coherence theorem of Etherington [Etherington, 1987]. A companion result is that the absence of ALL nonmonotonic loops (i.e., odd loops and non-zero even loops) guarantees a *unique* (though possibly inconsistent, as we have seen) well-founded labelling. This raises an interesting possibility. If we follow the approach described above for representing defaults by justifications of the form $\text{out}(\sim D) \rightarrow D$, and if only the contradiction-handler is allowed to produce justifications for the $\sim D$ propositions, then ordinary conflicts between defaults will only cause inconsistencies, not multiple (well-founded) labellings. Moreover, if the contradiction-handler is careful to avoid creating cycles in the nonmonotonic support structure, uniqueness of the labelling can be maintained. The guarantee of a unique extension/labelling would seem to be a nice property for a nonmonotonic reasoning theory.

VI. Contradiction Handling

Truth maintenance systems have been used to support a more efficient search process in problem solving. In this approach, the choices available in the search are represented as assumptions. An inconsistent set of choices gives rise to a contradiction, causing *dependency-directed backtracking* [Stallman and Sussman, 1977].

Planning applications combine temporal reasoning with problem solving search. If full use is to be made of a TMS in such an endeavor, then some assumptions will represent choices while others may represent default hypotheses about the environment (or even default persistences arising from the frame axioms).

Choices are generally ruled out if they conflict with our desires. Hypotheses are revised if they conflict with observation. When the two are mixed, trouble can result. Suppose, for example, our old friend Tweety is incarcerated in a bird cage, and we are considering opening the cage door. Under normal conditions, we can deduce that Tweety might fly away.

Let us suppose further that we do not wish Tweety to fly away. This conflict between an expectation and a desire would ordinarily lead to dependency-directed backtracking. However, if assumptions are represented uniformly, then a TMS could just as easily revise a default hypothesis about Tweety (say, that he has wings), as revise the choice of opening the birdcage. Thus, the system might postulate Tweety is wingless solely to avoid the disagreeable conclusion that he might fly away. That would be wishful thinking!

On the other hand, suppose we actually do want Tweety to fly away and open the birdcage for that purpose. In this case, if we wait patiently but observe no flight we might be justified in concluding that Tweety is abnormal in some way that is preventing the flight.

In other words, a conflict between an expectation and a desire leads one to reconsider choice of action, while a conflict between an expectation and an observation should lead one to reconsider one's beliefs. The TMS contradiction handling machinery will need to make such distinctions when employing *reductio ad absurdum* reasoning. This suggests that contradiction-handling needs to be treated as an extra-logical operation, rather than being built in to the logical formalism.

VII. Conclusions

Hanks and McDermott pointed out a difficulty in the default logic formalization of temporal reasoning: the existence of anomalous extensions. Closer examination shows the difficulty is not peculiar to temporal reasoning, but occurs in a wide range of default reasoning tasks.

In these cases, the natural representation of the problem in a truth maintenance system appears to clear up the difficulty. An inspection of the TMS representation suggests that nonnormal default rules are required to approximate its behavior. Further investigation indicates that nonnormal defaults (or some equivalent) are crucial in avoiding anomalous extensions.

We have also seen that nonnormal defaults can be isolated and represented by simple nonnormal default rules, or (using a TMS) nonmonotonic justifications of a simple form. An approach is suggested where conflicts between defaults cause inconsistencies rather than multiple extensions. In this approach the responsibility for resolving inconsistencies is shifted to an external contradiction-handler.

In support of the view that contradiction-handling should be regarded as an extra-logical operation, a new difficulty has been noted concerning the use of *reductio ad absurdum* reasoning in applications which combine default reasoning with problem solving. It appears that such reasoning needs to make distinctions -- between choices and hypotheses, and desires and observations -- which are at a level beyond ordinary logic.

Acknowledgements

I am grateful to Richard Fikes, Johan de Kleer and, particularly, Bob Nado for conversations which helped to clarify many of the ideas presented in this paper. The referees made helpful comments for improving the presentation. Discussions at the Frame Problem Workshop prompted my consideration of the differences between the TMS and Default Logic in responding to new information.

References

[Charniak *et al.*, 1979] Charniak, E., Riesbeck, C., and D. McDermott. *Artificial Intelligence Programming.* L. E. Erlbaum, Baltimore, 1979.

[Doyle, 1979] Doyle, J. A Truth Maintenance System. *Artificial Intelligence* 12(3), 1979.

[Etherington, 1987] Etherington, D.W. Formalizing Nonmonotonic Reasoning Systems. *Artificial Intelligence* 31:41-85, 1987.

[Hanks and McDermott, 1986] Hanks, S., and D. McDermott. Default Reasoning, Nonmonotonic Logics, and the Frame Problem. In *Proceedings AAAI-86.* Philadelphia, 1986.

[McCarthy, 1980] McCarthy, John. Circumscription - A Form of Non-Monotonic Reasoning. *Artificial Intelligence* 13:27-39, 1980.

[McCarthy, 1986] McCarthy, John. Applications of Circumscription to Formalizing Common-Sense Knowledge. *Artificial Intelligence* 28:89-116, 1986.

[Reiter, 1980] Reiter, Raymond. A Logic for Default Reasoning. *Artificial Intelligence* 13:81-132, 1980.

[Reiter and Criscuolo, 1983] Reiter, R. and G. Criscuolo. Some Representational Issues in Default Reasoning. *Int. J. Comput. Math.* 9:1-13, 1983.

[Shoham, 1986] Shoham, Y. Chronological Ignorance: Time, Nonmonotonicity, Necessity and Causal Theories. In *Proceedings AAAI-86.* Philadelphia, 1986.

[Stallman and Sussman, 1977] Stallman, R.M., and G.J. Sussman. Forward Reasoning And Dependency-directed Backtracking In A System For Computer-Aided Circuit Analysis. *Artificial Intelligence* 9:135-196, 1977.