mumble

Investigations of Model-Preference Defaults

Mark Boddy
Honeywell Systems and Research Center
boddy@src.honeywell.com

Robert P. Goldman
Tulane University
rpg@cs.tulane.edu

Keiji Kanazawa
University of British Columbia
kanazawa@cs.ubc.ca

Lynn Andrea Stein
Massachusetts Institute of Technology
las@ai.mit.edu

Abstract

Selman and Kautz's award-winning paper on the logic of model-preference defaults introduced a new language and tractability results. This paper assesses the contributions of that work, including the properties of the logic itself, its relationship to some previous logics, and its applicability to some known tractable problems.

Submitted to *Fundamenta Informaticae*
Special Issue on Nonmonotonic Reasoning

1 Introduction

In an award-winning paper [11], Selman and Kautz presented the logic of *model-preference defaults* (MPD). MPD is a family of propositional default logics intended as a testbed for exploring nonmonotonic reasoning. Selman and Kautz define the search problem of finding a maximal model given an MPD axiomatization. They show that for general MPD axiomatizations this search problem is $\mathcal{NP}$-hard. They describe a polynomial search algorithm for an MPD language limited to acyclic Horn-form defaults with a monotonic theory containing only literals.

These results raise some further questions: What are the expressive limitations of MPD? What insights does MPD give us into other tractable formalisms? Defining maximal models in terms of defaults is intuitively appealing; do such models correspond to the models or extensions derivable using other logics?

Our paper is a further exploration of MPD. Although connections between that logic and several others are made here, it is not clear that these connections contribute to the understanding of the logics other than MPD; rather, they help us to place MPD in the context of much previous research. In the end, Selman and Kautz have given us a family of languages with identified tractability properties; however, these languages are not themselves generally useful, nor do they shed much light on the tractability of previous languages. In this sense, the current paper is something of an indictment of MPD.

2 Background and Terminology

This section briefly reviews the logic of MPD as presented by Kautz and Selman [11]. The reader familiar with their work may wish to skip to the next section.

An MPD *axiomatization* is a pair $\langle T, D\rangle$. T is a set of formulae in the propositional calculus, referred to as the (monotonic or propositional) *theory*. D is a set of defaults of the form $\alpha \rightarrow \beta$, where α is a conjunction of literals, and β is a single literal. Selman and Kautz use the set of defaults to define a partial order, the *preference relation*, over the models of the propositional theory. Intuitively, this ordering prefers models in which more default rules

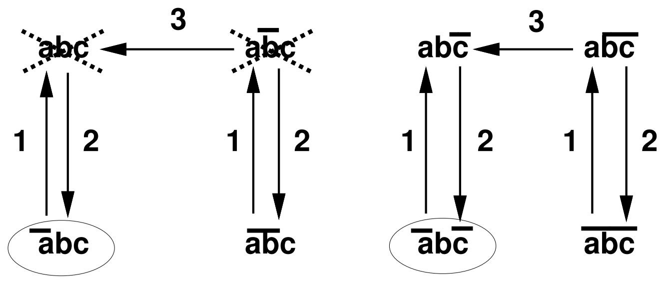

$D = \{1 : \rightarrow \mathsf{a},\ 2 : \rightarrow \bar{\mathsf{a}},\ 3 : \mathsf{a} \rightarrow \mathsf{b}\}$ $\qquad$ $T = \{\bar{\mathsf{a}} \vee \bar{\mathsf{c}}\}$

Figure 1: Preference relation graph

"fire."

The preference relation is defined in terms of an accessibility relation induced by the set of defaults. For any default rule $d : \alpha \rightarrow \beta$ and truth assignment $\mathcal{T}$ satisfying α, $d(\mathcal{T})$ is the truth assignment satisfying α and β, and otherwise identical to $\mathcal{T}$. If $\mathcal{T} \models \beta$, $d(\mathcal{T}) = \mathcal{T}$. MPD axiomatizations may make use of a *specificity relation*. A default $d' : \alpha \wedge \gamma \rightarrow \overline{\beta}$ blocks (is more specific than) d at $\mathcal{T}$ if $\mathcal{T} \models \alpha \wedge \gamma$. If d is not blocked at $\mathcal{T}$, then $d(\mathcal{T})$ is *accessible* from $\mathcal{T}$. The preference relation is the reflexive transitive closure of accessibility. One model is strictly preferred to another if the former is preferable to the latter and not vice versa. A model is *maximal* if there is no model strictly preferred to it.

Figure 1 shows the accessibility relation on truth assignments for the MPD axiomatization $D = \{1 : \rightarrow \mathsf{a},\ 2 : \rightarrow \bar{\mathsf{a}},\ 3 : \mathsf{a} \rightarrow \mathsf{b}\}$, $T = \{\bar{\mathsf{a}} \vee \bar{\mathsf{c}}\}$. The numbers on the edges indicate the corresponding default rules. The circled truth assignments are maximal models; truth assignments not satisfying T are crossed out. Note that the accessibility relation goes through these non-models, so that, e.g., $\overline{\mathsf{ab}}\mathsf{c}$ is not maximal, because $\bar{\mathsf{a}}\mathsf{bc}$ is both more preferable and a model of $\mathcal{T}$.

Selman and Kautz consider restricted MPD languages containing *acyclic* sets of defaults. For any set of defaults, we can define a graph with vertices corresponding to the set of propositional literals. For each default $\alpha \rightarrow \beta$, for each propositional letter α_i mentioned in α, add a directed arc $\langle \alpha_i, \beta \rangle$. A set of defaults is acyclic if its corresponding graph is acyclic.

Selman and Kautz define a number of sublanguages of MPD and analyze the corresponding search problems. The search problem for an MPD language is the problem of finding any single, maximal model for a given axiomatization. $\mathcal{D}$ is the unrestricted MPD language without specificity. $\mathcal{D}^+$ is $\mathcal{D}$ with the addition of the specificity relation. $\mathcal{D}^+_{\mathrm{a}}$ restricts the defaults of $\mathcal{D}^+$ to be acyclic. A language with the specificity relation is marked with a $+$ superscript. Acyclic languages are marked with a subscripted $_{\mathrm{a}}$. $\mathcal{DH}$ is the set of MPD axiomatizations with "Horn-form" defaults: if $\alpha \to \beta$ is a default, then α contains no negated literals. $\mathcal{DH}^+$ adds specificity; $\mathcal{DH}_{\mathrm{a}}$ and $\mathcal{DH}^+_{\mathrm{a}}$ require that the defaults be acyclic and Horn-form.

In [12], Selman and Kautz show that the search problems for $\mathcal{D}$, $\mathcal{D}^+$ and $\mathcal{DH}^+$ with empty theories are $\mathcal{NP}$-hard. There is a linear algorithm for $\mathcal{DH}$ with an empty propositional theory, but when $\mathcal{DH}$ is augmented by a Horn-form monotonic theory, it becomes $\mathcal{NP}$-hard. Selman and Kautz describe a polynomial algorithm for $\mathcal{DH}^+_{\mathrm{a}}$ with a monotonic theory made up of literals. An anonymous reviewer has also demonstrated that $\mathcal{D}^+_{\mathrm{a}}$ is $\mathcal{NP}$-hard.

3 Expressive Power

The expressiveness of model-preference defaults is limited by two features of the language. First, "strengthening" an axiomatization by replacing a default with the corresponding monotonic (material) implication can actually lead to a weaker theory. Second, when the classical part of the axiomatization is limited to a set of literals (i.e. without potentially anomalous material implications), model-preference default axiomatizations are incapable of properly expressing situations involving negation.

3.1 Propositional Theories

Selman and Kautz's criteria for model maximality leads to anomalous results in the presence of monotonic Horn theories. Attempts to strengthen an axiomatization by replacing a default with the corresponding monotonic implication may actually result in a *weaker* theory. For example, consider drawing objects from a bag containing mostly balls. We might axiomatize this as $\mathrm{D} = \{1 : \to \mathsf{ball},\ 2 : \mathsf{ball} \to \ \mathsf{round}\}$; $\mathrm{W} = \{\}$. This has one maximal model in which we expect to pull out a round ball. If we exclude footballs

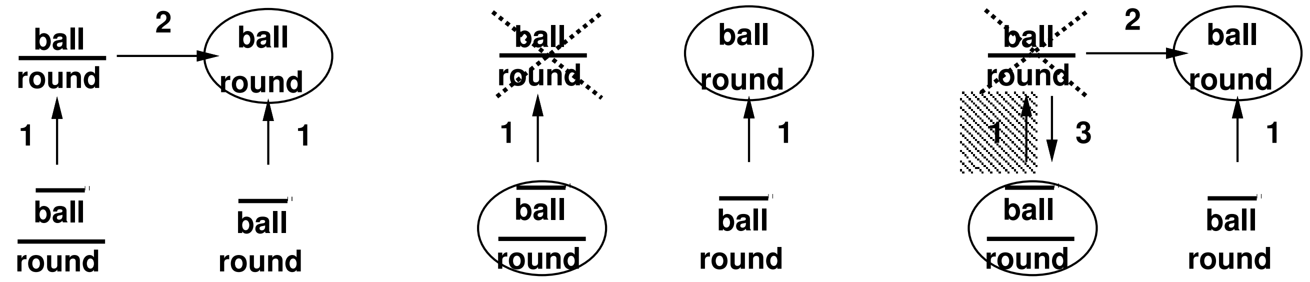

(i)

$$D = \{\ 1: \ \rightarrow \mathsf{ball},\quad 2: \mathsf{ball} \rightarrow \mathsf{round}\ \}$$

(ii)

$$D = \{\ 1: \ \rightarrow \mathsf{ball}\ \}$$
$$T = \{\ \mathsf{ball} \supset \mathsf{round}\ \}$$

(iii)

$$D = \{\ 1: \ \rightarrow \mathsf{ball},\quad 2: \mathsf{ball} \rightarrow \mathsf{round},\quad 3: \overline{\mathsf{round}} \rightarrow \overline{\mathsf{ball}}\ \}$$
$$T = \{\ \mathsf{ball} \supset \mathsf{round}\ \}$$

Figure 2: Accessibility graphs

from the bag, we can strengthen our assumptions: D $= \{\ 1 : \ \rightarrow\ \mathsf{ball}\}$; W $= \{\mathsf{ball} \supset \mathsf{round}\}$. However, this has the unexpected consequence that there are now *two* maximal models, in only one of which we draw a round ball. The accessibility graphs for these axiomatizations are shown in figures 2 i. and 2 ii., respectively. This has the counter-intuitive result that a seemingly stronger axiomatization has weaker consequences.

We could try to fix this by adding the default rule corresponding to any monotonic implication. Unfortunately, there is in general no way to distinguish which of several defaults to add: given $\mathsf{a} \supset \overline{\mathsf{b}} = \overline{\mathsf{a}} \vee \overline{\mathsf{b}}$, we could add either $\mathsf{a} \rightarrow \overline{\mathsf{b}}$ or $\mathsf{b} \rightarrow \overline{\mathsf{a}}$. Adding all of the corresponding default rules may result in the same anomalous maximal models as not adding any of them. Figure 2 iii. shows the result of adding both defaults corresponding to $\mathsf{ball} \supset \mathsf{round}$.

This is a direct consequence of the contraposability of $\supset$, as opposed to the directionality of MPD defaults. In those cases where the material implication has direction—as, for example, in representing strict inheritance (see section 5.1, below)—the strategy of adding appropriate defaults may be more successful.

3.2 Negation in Tractable MPD Sublanguages

One of the weaknesses of the tractable MPD language $\mathcal{DH}_a^+$—acyclic Horn-form defaults with specificity—is the expressive poverty of Horn-form defaults. In particular, prohibiting negated literals in the antecedents of default rules makes it impossible to reason by cases. Consider representing a situation in which shooting at someone when the light is on typically kills that person, while shooting at someone in the dark typically kills an innocent bystander.

In representing this situation, we might use dual propositions, dark and light. We could then write rules $\mathsf{shoot} \wedge \mathsf{light} \rightarrow \mathsf{kill\_target}$ and $\mathsf{shoot} \wedge \mathsf{dark} \rightarrow \mathsf{kill\_bystander}$. Unfortunately, there is no satisfactory way to relate the truth values of light and dark. Defining dark as logically equivalent to $\overline{\mathsf{light}}$ requires the ability to express biconditionals. This cannot be accomplished using acyclic defaults alone; such defaults will not rule out a model in which light and dark coexist. Nor is it possible to express the required biconditional using cyclic but Horn-form propositional theories. Indeed, even using (intractable) $\mathcal{DH}^+$ defaults, such aliasing requires that dark *always* be used in place of $\overline{\mathsf{light}}$, and light in place of $\overline{\mathsf{dark}}$.

4 Vivid reasoning

The MPD search problem involves finding a single arbitrary maximal model. Selman and Kautz's tractability results do not address the problems of goal-directed reasoning—Is there a maximal model endorsing some conclusion?—or skeptical reasoning—Do all maximal models endorse that conclusion? In other work, Kautz and Selman [5] suggest that the computational complexity of these problems is likely to be significantly worse than the arbitrary extension (model-finding) problem. This section explores the ramifications of the restriction to model-finding for vivid reasoning. Similar issues arise in sections 5 and 6, below.

Selman and Kautz offer MPD as a logic suited for vivid reasoning. Levesque [6] defines vivid knowledge bases as databases representing all facts of interest in an explicit way.[1] Vivid knowledge bases allow us to substitute simple database operations for more costly forms of reasoning. Brachman and Etherington [1] suggest that since vivid reasoning will not be able to capture the full range of problem-solving behavior, an intelligent program should have both a vivid knowledge base and a general problem-solver for more difficult "puzzle-mode" reasoning.

The MPD search algorithm is to return a single maximal model of an axiomatization, regardless of whether there is only one such model or there are many. Constructing such a model may require making arbitrary truth assignments to some literals. Sometimes, such choices are inappropriate, as for example in Reiter and Criscuolo's "Nixon diamond" [10]. The construction of a complete truth assignment requires this overcommitment, analogous to McDermott and Doyle's *brave* or Touretzky et al.'s *credulous* reasoners [8, 17].

Since the tractability of MPD depends on the construction of a single maximal model, we cannot quantify over maximal models to determine, whether a particular literal has the same truth value in all such models or in any such models. Inference from a single MPD maximal model thus blurs the distinction between truth, likelihood, and arbitrary choice. Since it cannot be determined from an MPD model which truth values are the result of arbitrary choices, MPD is not suited for use in a hybrid reasoning system like that proposed by Etherington and Brachman. Such a system would be unable to determine when to refer a query to its more accurate—but more expensive—"puzzle-mode" problem solver. An MPD model may represent one "vivification" of a knowledge base, but it provides no focus or support for further, more careful, inference.

5 Applications to other formalisms

In this section, we discuss MPD's suitability as a formalism for various known tractable problems. MPDs with Horn-form theories can be used to represent inheritance hierarchies with strict and defeasible links, yielding a polynomial time algorithm for inheritance comparable to that of Stein [15]. However, the addition of specificity relies on an elaborate translation mechanism. In any case, the result differs from standard inheritance theories in that it is

[1] Explicit is taken to mean directly represented or represented as a simple relation among expressions.

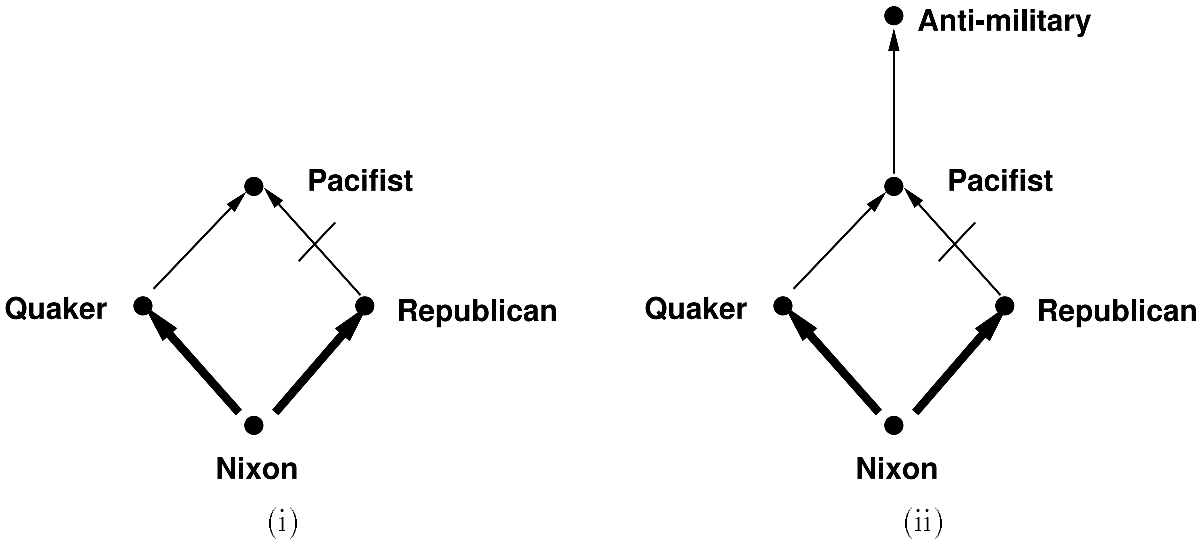

Figure 3: Some simple inheritance hierarchies

arbitrary and overly determined. Temporal projection for propositional theories and totally ordered events —in particular, theories of the form defined by Shoham in his logic of chronological ignorance—can be axiomatized in $\mathcal{D}^+_{\mathrm{a}}$ or $\mathcal{D}^+$. However, the complete and arbitrary nature of MPD again leads to a slightly different conclusion set.

5.1 Inheritance

Kautz and Selman say that MPD—in particular, $\mathcal{DH}^+_{\mathrm{a}}$—"can encode inheritance hierarchies of the form examined by Touretzky, but are strictly more general." They do not, however, provide any description of how this translation is to be accomplished. In this section, we explore one such translation and examine the adequacy of MPD for inheritance.

Touretzky et al. [17] identify several dimensions along which existing theories of inheritance vary. The version of inheritance originally proposed by Touretzky [16] is downwards—it involves reasoning about which objects have a given property or belong to a given class. In contrast, the inheritance reasoning of MPD is upwards—it concerns the properties possessed by a given entity. The differences between these approaches are detailed elsewhere [13, 17]; most notably, Levesque and Selman demonstrate that most upwards reasoning is tractable while downward reasoning is, in general, $\mathcal{NP}$-hard.

To translate an inheritance hierarchy into MPD, we can simply replace the edges of the hierarchy—e.g., that Quakers are generally Pacifists—with the corresponding default rule: $\mathsf{Q} \rightarrow \mathsf{P}$. Strict links—such as "Nixon *is* a Quaker"—are replaced by the corresponding default—$\mathsf{N} \rightarrow \mathsf{Q}$—*and* the material implication—$\mathsf{N} \supset \mathsf{Q}$. Note that, while it is not in general possible to determine the "direction" of the material conditional—$\mathsf{N} \supset \mathsf{Q}$ versus $\overline{\mathsf{Q}} \supset \overline{\mathsf{N}}$—and hence to derive only the appropriate default—$\mathsf{N} \rightarrow \mathsf{Q}$, rather than $\overline{\mathsf{Q}} \rightarrow \overline{\mathsf{N}}$—in the case of inheritance the hierarchy provides the disambiguating information.

The translation of the hierarchy in figure 3 i. (with strict links represented as wide arrows and defeasible ones as narrow arrows) is then given by

$$T = \left\{ \begin{array}{c} \mathsf{N} \supset \mathsf{Q} \\ \mathsf{N} \supset \mathsf{R} \end{array} \right\} \qquad D = \left\{ \begin{array}{c} \mathsf{N} \rightarrow \mathsf{Q} \\ \mathsf{N} \rightarrow \mathsf{R} \\ \mathsf{Q} \rightarrow \mathsf{P} \\ \mathsf{R} \rightarrow \overline{\mathsf{P}} \end{array} \right\}$$

If we add N to T—i.e. if we assert "Nixon"—the resulting MPD axiomatization has two preferred models: NQRP, in which Nixon is a Pacifist, and $\mathsf{NQR}\overline{\mathsf{P}}$, in which he is not.[2] This corresponds to the results of standard inheritance reasoners. In particular, it is *credulous* reasoning—it allows either of the two possibilities. In contrast, skeptical reasoning supports neither of contradictory conclusions.

There are differences between MPD and standard inheritance as well. Consider figure 3 ii., in which we've added the default that pacifists are generally anti-military. The resulting MPD axiomatization has three preferred models: NQRPA, $\mathsf{NQR}\overline{\mathsf{P}}\mathsf{A}$, and $\mathsf{NQR}\overline{\mathsf{PA}}$. In contrast, standard inheritance reasoners yield only two interpretations of the hierarchy: one corresponding to NQRPA, and one corresponding to the *partial* model $\mathsf{NQR}\overline{\mathsf{P}}$, with A uncertain. That is, if Nixon is not a Pacifist, it is completely unknown—not merely undecided—whether Nixon is Anti-military. Because MPD assigns true or false to *every* literal, this "unknown" state cannot be captured.

[2] Note, however, that Selman and Kautz's polynomial time algorithm for $\mathcal{DH}^{+}_{\mathrm{a}}$ will find

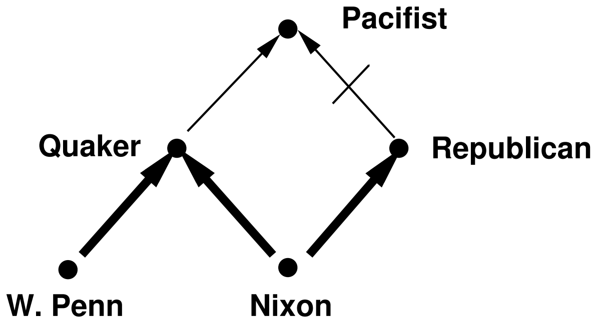

$$\begin{array}{rlcl} D = & \{\mathsf{W} & \rightarrow & \mathsf{Q} \\ & \mathsf{N} & \rightarrow & \mathsf{Q} \\ & \mathsf{N} & \rightarrow & \mathsf{R} \\ & \mathsf{Q} & \rightarrow & \mathsf{P} \\ & \mathsf{R} & \rightarrow & \overline{\mathsf{P}}\} \\ T = & \{\mathsf{W} & \supset & \mathsf{Q} \\ & \mathsf{N} & \supset & \mathsf{Q} \\ & \mathsf{N} & \supset & \mathsf{R}\} \end{array}$$

Figure 4: Anomalies arise when reasoning about Penn, i.e., when W is added to T.

The problem is exacerbated when we reason about hierarchies with multiple leaves or about non-leaf nodes. For example, figure 4 includes an additional additional leaf node: William Penn, prototypical Quaker. If we reason about Penn—i.e., we add W to T—we are forced to conclude either N or $\overline{\mathsf{N}}$. But the consequences are even more severe here: if we conclude N, we must also adopt R (since $\mathsf{N} \supset \mathsf{R}$); and now we have reason to conclude $\overline{\mathsf{P}}$ (by $\mathsf{R} \rightarrow \overline{\mathsf{P}}$). Thus, the MPD translation of figure 4 together with W allows five maximal models: WNQRP, $\mathsf{WNQR\overline{P}}$, $\mathsf{W\overline{N}Q\overline{R}P}$, $\mathsf{W\overline{N}QR\overline{P}}$, and $\mathsf{W\overline{N}QRP}$. While Selman and Kautz's polynomial-time procedure will find only one of these, MPD itself supports all and hence fails to validate the conclusion that William Penn is a Pacifist.

While the translation from inheritance to MPD is awkward for hierarchies like those of figures 3 and 4, it is far more cumbersome for inheritance theories that make use of some form of preemption. For example, in figure 5, information about penguins is more specific to Tweety than information about birds, so we'd like to conclude that Tweety doesn't fly. To accomplish this, we need to invoke a preemption detection scheme. Here, we show how to adapt that of Stein [15].

First, for each node, x, we create a unique new propositional variable—

---

only one of these models. (See section 4, above.)

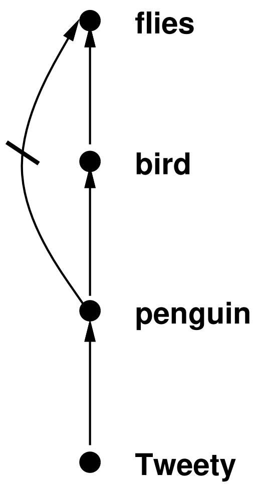

Figure 5: An inheritance hierarchy with preemption

say, $\hat{x}$. We will use $\hat{x}$ to indicate that x is the node whose properties we're reasoning about. For example, we'd add $\hat{\mathsf{N}}$ to the propositional theory in the examples of the previous paragraph. By default, $\hat{x}$ will be set false—$\rightarrow \overline{\hat{x}}$—and whenever $\hat{x}$ is set true in the propositional theory—whenever we're reasoning about x—we'll want x to be true as well—$\hat{x} \rightarrow x$, and $\hat{x} \supset x$, for all x.

Now, whenever specificity overrides an inference—e.g., Tweety doesn't fly, because penguin is more specific than bird—we add a new default rule—in this case, $\hat{\mathsf{T}} \wedge \mathsf{P} \wedge \mathsf{B} \rightarrow \overline{\mathsf{F}}$ (when reasoning about Tweety flying, penguin overrides bird). MPD's rule of specificity now overrides the default rule $\mathsf{B} \rightarrow \mathsf{F}$.

The complete translation procedure is given in figure 6. Its complexity is polynomial—$O(n^6)$.[4] It yields an MPD axiomatization that resembles one particular flavor of inheritance, although it is still subject to the anomalies mentioned above. The resulting MPD inheritance reasoner is upwards—it reasons about the properties of a given node—and credulous—it always commits to a single, complete world-view. It uses a simple ("on-path") pre-

[3] x is reachable from a if there are a sequence of positive edges, or a sequence of positive edges followed by a *single* negative edge, from a to x.

[4] For a detailed analysis of complexity, see [15, Theorem 3], from which this algorithm is derived.

Inheritance-to-MPD

Let Γ be an inheritance hierarchy.
For each node a
  Let Γ_a be Γ
  For each node x reachable from a,[3] in topological order ; from a to ...
    For each (positive or negative) edge from p to x, in reverse topological order
      If the edge from p to x is positive,
        Let $\Gamma_a^\star = \Gamma_a - \{q_i \mid \text{there is a negative edge from } q_i \text{ to } x \text{ in } \Gamma_a\}$
        If there is no longer a sequence of positive edges from a to p in $\Gamma_a^\star$, then
          • remove the (positive) edge from p to x from Γ_a
          • add the default rule $\hat{a} \wedge \bigwedge_i q_i \wedge p \rightarrow \overline{x}$ to D
      Else if the edge from p to x is negative,
        Let $\Gamma_a^\star = \Gamma_a - \{q_i \mid \text{there is a positive edge from } q_i \text{ to } x \text{ in } \Gamma_a\}$
        If there is no longer a sequence of positive edges from a to p in $\Gamma_a^\star$, then
          • remove the (negative) edge from p to x from Γ_a
          • add the default rule $\hat{a} \wedge \bigwedge_i q_i \wedge p \rightarrow x$ to D
    For each remaining positive edge from p to x
      If Γ_a contains a sequence of positive edges from p through some nodes
        $r_1, \ldots, r_n$ to x, $n \geq 1$, such that there is no sequence of positive edges
        followed by a single negative edge from a to any of r_i, x,
          then remove the edge $p \cdot x$ from Γa
    If there is no longer a sequence of positive edges from a to x,
      then remove x from Γ_a.

Figure 6: Translation procedure for inheritance hierarchies

emption strategy. Unlike all previous inheritance reasoners, MPD requires the assignment of truth-values to *all* literals; many other systems exploit the locality of inference an “unknown” truth value allows. Further, much of the “work” of inheritance is captured by the translation procedure rather than the MPD’s own specificity relation. In general, it is not likely that MPD will be suited to simulation of Touretzky-style inheritance, or to capturing the more complicated (“off-path”) versions of preemption found in much of the literature.

5.2 Chronological Minimization

Shoham [14] describes a model of temporal reasoning based on the logic of *chronological maximal ignorance* (CMI). CMI uses the idea that time moves forward to solve two major problems of reasoning about action and change over time: qualification and extended prediction. In this section, we present a translation from Shoham's CMI framework to model preference defaults in $\mathcal{D}^{+}_{(\mathrm{a})}$.

Shoham begins with sentences of the form

$$\bigwedge_{i=1}^{m} \Box\varphi_i \wedge \bigwedge_{j=1}^{n} \Diamond\overline{\psi_j} \supset \Box\omega$$

where the times of φ_i and ψ_i are earlier than the time of ω. Intuitively, this rule means that if the preconditions hold ($\bigwedge_{i=1}^{m} \Box\varphi$), and we can reasonably assume that the qualifications do not interfere ($\bigwedge_{j=1}^{n} \Diamond\overline{\psi_j}$), then we conclude the consequent ($\Box\omega$).

To translate from CMI to MPD, we replace each CMI sentence with a set of defaults:

$$\begin{aligned} & & \textstyle\bigwedge_{i=1}^{m} \varphi_i \rightarrow \omega & \qquad (1)\\ \text{For each } \Diamond\psi_j, & & \textstyle\bigwedge_{i=1}^{m} \varphi_i \wedge \psi_j \rightarrow \overline{\omega} & \qquad (2)\\ & \text{and} & \textstyle\bigwedge_{i=1}^{m} \varphi_i \wedge \psi_j \rightarrow \omega & \qquad (3)\\ \text{For each } \Diamond\psi_j, & & \rightarrow \overline{\psi_j} & \qquad (4) \end{aligned}$$

Default rule (1) says that if the preconditions hold, we assume the consequent. (2) and (3) allow the qualifications to interfere: if any one of the qualifications holds, we can't assume the consequent; i.e., the consequent and its negation become equipreferable in the absence of further information. In this case, (1) will be blocked by the more specific (2) and (3). Finally, (4) ensures that the qualifications do not, in general, hold. (4) will apply unless something *forces* the qualification (ψ_j) to hold.

We can allow states to persist by modifying the set of defaults slightly. We use the specificity criterion to allow change-rules to override persistences. If $\omega = (t, p)$, and p is a *state*, we define $\omega_{-1} = (t-1, p)$. For each such state, we add a persistence rule:

$$\omega_{-1} \rightarrow \omega \qquad (5)$$

From CMI to MPD

Given a CMI theory containing:

$$\Box(t, \mathsf{turn\_key}) \wedge \Diamond\neg(t, \mathsf{potato\_in\_tailpipe}) \supset \Box(t+1, \mathsf{started\_motor})$$
$$\Box(t, \mathsf{put\_potato\_in\_tailpipe}) \supset \Box(t+1, \mathsf{potato\_in\_tailpipe})$$

we construct an MPD axiomatization which includes the schemata:

$$(t, \mathsf{turn\_key}) \rightarrow (t+1, \mathsf{started\_motor}) \qquad \text{by (1)}$$

$$(t, \mathsf{turn\_key}) \wedge (t, \mathsf{potato\_in\_tailpipe}) \rightarrow \neg(t+1, \mathsf{started\_motor}) \qquad \text{by (2)}$$
$$(t, \mathsf{turn\_key}) \wedge (t, \mathsf{potato\_in\_tailpipe}) \rightarrow (t+1, \mathsf{started\_motor}) \qquad \text{by (3)}$$

$$\neg(t, \mathsf{started\_motor}) \wedge (t, \mathsf{turn\_key}) \rightarrow (t+1, \mathsf{started\_motor}) \qquad \text{by (6)}$$

$$(t, \mathsf{started\_motor}) \wedge$$
$$\quad (t, \mathsf{turn\_key}) \wedge (t, \mathsf{potato\_in\_tailpipe}) \rightarrow \neg(t+1, \mathsf{started\_motor}) \qquad \text{by (7)}$$
$$(t, \mathsf{started\_motor}) \wedge$$
$$\quad (t, \mathsf{turn\_key}) \wedge (t, \mathsf{potato\_in\_tailpipe}) \rightarrow (t+1, \mathsf{started\_motor}) \qquad \text{by (8)}$$

$$\neg(t, \mathsf{potato\_in\_tailpipe}) \qquad \text{by (4)}$$

$$(t, \mathsf{put\_potato\_in\_tailpipe}) \rightarrow (t+1, \mathsf{potato\_in\_tailpipe}) \qquad \text{by (1)}$$
$$\neg(t, \mathsf{potato\_in\_tailpipe}) \wedge$$
$$\quad (t, \mathsf{put\_potato\_in\_tailpipe}) \rightarrow (t+1, \mathsf{potato\_in\_tailpipe}) \qquad \text{by (6)}$$

$$(t, \mathsf{potato\_in\_tailpipe}) \rightarrow (t+1, \mathsf{potato\_in\_tailpipe})$$
$$\neg(t, \mathsf{potato\_in\_tailpipe}) \rightarrow \neg(t+1, \mathsf{potato\_in\_tailpipe})$$
$$(t, \mathsf{started\_motor}) \rightarrow (t+1, \mathsf{started\_motor}) \qquad \text{by (5)}$$
$$\neg(t, \mathsf{started\_motor}) \rightarrow \neg(t+1, \mathsf{started\_motor})$$

and

$$\rightarrow \neg(t, \mathsf{put\_potato\_in\_tailpipe})$$
$$\rightarrow \neg(t, \mathsf{turn\_key}) \qquad \text{by (9)}$$

Figure 7: A sample CMI translation

We now augment (1), (2) and (3) with modified rules which allow them to override these persistences.[5] We do this by adding $\overline{\omega}_{-1}$ to the left-hand sides:

$$\overline{\omega}_{-1} \wedge \bigwedge_{i=1}^{m} \varphi_i \rightarrow \omega \tag{6}$$

and

$$\text{For each } \Diamond\psi_j, \quad \overline{\omega}_{-1} \wedge \bigwedge_{i=1}^{m} \varphi_i \wedge \psi_j \rightarrow \overline{\omega} \tag{7}$$
$$\text{and } \Diamond\psi_j, \quad \overline{\omega}_{-1} \wedge \bigwedge_{i=1}^{m} \varphi_i \wedge \psi_j \rightarrow \omega \tag{8}$$

Finally, in any case, for each $\varphi = (t, p)$, p an *action*, we need a default:

$$\rightarrow \overline{\varphi} \tag{9}$$

Unlike states, actions and their negations are asymmetric. Unless we have reason to believe otherwise, *actions do not occur*. An example of the translation of a CMI theory is given in figure 7.

The MPD translation of a CMI theory corresponds to the original, but it is not identical. The problem is that MPD models are fully determined—every proposition is assigned either *true* or *false*—while CMI uses the epistemic operator ($\Box$) to create a third truth value. In general, each of the MPD models will correspond to one of the possible worlds over which CMI's epistemic operator implicitly quantifies, and the CMI model will be directly described by the intersection of MPD's preferred models. However, as we have seen in section 4, above, MPD model-finding algorithms identify only a single maximal model. Further, the translation that we have given here falls into $\mathcal{D}^+$ or $\mathcal{D}^+_{\mathrm{a}}$, both MPD languages for which even model finding is known to be $\mathcal{NP}$-hard. Thus, while we have captured roughly the same intuitions as Shoham's CMI formalization, we have sacrificed the efficiency of his special purpose inference method.

5.3 The parking lot problem

One difficulty with Shoham's CMI was raised by Kautz [4]. CMI delays action as long as possible, forcing us to conclude that actions happen only at the last possible moment. In his parking-lot scenario, Kautz suggests that

[5] We retain the old versions—(1)–(3)—to apply in case ω_{-1} may be true.

knowing a car has been stolen should *not* lead to the belief that it was stolen immediately before the discovery of the theft.

Kautz poses the following problem: Suppose that you leave your car in a parking lot, and come back later to find it gone. Suppose further that you know that the act of stealing a car is the only possible cause for its disappearance. How might this be axiomatized so that we can conclude that the car was, in fact, stolen? There are several nonmonotonic logics within which this sort of reasoning can be done (e.g. [7, 8, 9]). As Hanks and McDermott [2] point out, these logics all suffer from the difficulty that there will in general be more than one valid extension from a given theory, where only some of those extensions correspond to conclusions that we intuitively see as correct.

Various fixes have been proposed for this problem. In formal systems for temporal reasoning, a popular solution has been to incorporate the asymmetric structure of time into the process of selecting extensions to a given theory [4, 14]. Informally, the effect of these approaches might be stated as "as little happens as late as possible." The problem with this assumption for the current example is that there is no reason to prefer an extension in which the car is stolen at the last possible moment over one in which it was stolen earlier. In this case, there should be multiple extensions, in each of which the car is stolen at a different time.

This example can be axiomatized using the following MPD schemata:

- $\mathsf{car}_t \rightarrow \mathsf{car}_{t+1}$; $\overline{\mathsf{car}}_t \rightarrow \overline{\mathsf{car}}_{t+1}$ – forward persistence.
- $\mathsf{car}_t \rightarrow \mathsf{car}_{t-1}$; $\overline{\mathsf{car}}_t \rightarrow \overline{\mathsf{car}}_{t-1}$ – backward persistence.
- $\rightarrow \overline{\mathsf{steal}}_t$ – minimization of actions.
- $\mathsf{car}_t \wedge \mathsf{steal}_t \rightarrow \overline{\mathsf{car}}_{t+1}$ – causation.
- $\overline{\mathsf{car}}_{t+1} \wedge \mathsf{car}_t \rightarrow \mathsf{steal}_t$ – explanation.

Suppose that the only time points that our universe contains are 1, 2, and 3. There are two maximal models for this set of defaults in the absence of any theory, corresponding to the car being either present for all three points or absent for all three, and no theft having taken place.

Now consider the theory $\{\mathsf{car}_1, \overline{\mathsf{car}}_3\}$, corresponding to observations made at the time points 1 and 3. There are two maximal models for this theory, corresponding to the car's getting stolen at time 1 or time 2. This is, in fact, what is desired—either model is acceptable. The need for both forward and backward persistence results in a set of cyclic defaults, and thus an axiomatization that is not contained in any of the known polynomial sublanguages of MPD.

Using the specificity relation to block defaults overrides persistences in the appropriate manner. This eliminates the need to make strong assumptions about a "direction" in time in order to establish the precedence of persistence over spontaneous change, and of causation over persistence.

Starting with the theory $\{\mathsf{car}_1, \mathsf{steal}_2\}$ we want one maximal model, in which the car is present exactly until it is stolen. There are axioms we might add to the set above, corresponding to stealing the car causing it to be gone, and to the car being gone not persisting back through the theft. Unfortunately, this will create a cycle, resulting in a set of maximal models, rather than the single model we would prefer. However, we can add the axiom $\mathsf{car}_t \wedge \mathsf{steal}_t \rightarrow \overline{\mathsf{car}}_{t+1}$ and remove backward persistence, and get the model we desire. (This is an axiomatization corresponding to a CMI theory.)

6 Comparisons to other logics

In this section, we investigate the relation between MPD and other formalisms for nonmonotonic reasoning. The style of reasoning embodied in circumscription is consistent with MPD, but that of default logic is based on a fundamentally different model of default reasoning.

6.1 Circumscription

Circumscription results in models that are minimal in the satisfaction of some formulae. MPD is concerned with finding maximal models, where the ordering relation is derived from a set of defaults. Intuitively, it seems that finding the correct way to encode the preference relation will allow an exact translation from MPD to circumscription. In fact, the more obvious translation from MPD axiomatizations to circumscriptive theories does not work.

In circumscriptive theories, default rules are typically expressed using abnormality predicates. Thus, we might translate the MPD default $\alpha \rightarrow \beta$ as $\alpha \wedge \overline{\mathsf{ab}_1} \supset \beta$ (α implies β, unless there is something abnormal about the situation). Unfortunately, the two forms are not equivalent: the circumscriptive

MPD-to-Circumscription

1. Rewrite the defaults to include specificity explicitly.
2. Given the set of literals $\mathcal{L} = \{l_1, l_2, \ldots, l_n\}$, for each truth assignment $\mathcal{T}$, let $[\mathcal{T}] = \bigwedge_i \lambda_i$, where $\lambda_i = l_i$ if $\mathcal{T} \models l_i$ and $\overline{l}_i$ otherwise. Also, for any formula $\mathcal{F}$, let $[\![\mathcal{F}]\!] = \circ_i \lambda_i$, where $\circ$ is alphabetically ordered string concatenation.
3. For each truth assignment $\mathcal{T}$, add the axiom $[\mathcal{T}] \supset \mathsf{ab}([\![[\mathcal{T}]]\!])$.
4. For each default of the form $\alpha \to \beta$, add the axioms corresponding to the following schema:

$$\mathsf{ab}([\![\alpha \wedge \overline{\beta} \wedge \mathsf{X}]\!]) \supset \mathsf{ab}([\![\alpha \wedge \beta \wedge \mathsf{X}]\!])$$

   where X varies over the 2^n possible n-element satisfiable conjunctions of the (negated and unnegated) literals not appearing in $\alpha \wedge \beta$.
5. Add as axioms the propositional theory.
6. Circumscribe $\mathsf{ab}(x)$ over the resulting theory.

Figure 8: Translation procedure to circumscription

rule contraposes, while the MPD axiom does not. The default rule $\alpha \to \overline{\beta}$ would translate as $\alpha \wedge \overline{\mathsf{ab}_1} \supset \overline{\beta}$. This formula is equivalent to $\beta \wedge \overline{\mathsf{ab}_1} \supset \overline{\alpha}$ and would thus also serve as a translation of the MPD default $\beta \to \overline{\alpha}$. The two default rules are not equivalent in MPD, and so this translation is incorrect. Even where contraposition poses no difficulty, the circumscriptively minimal models may differ from the maximally preferable models under MPD. For example, given the default $\alpha \to \beta$ and the propositional theory $\overline{\alpha} \vee \overline{\beta}$, there are three maximal models: $\alpha\overline{\beta}$, $\overline{\alpha}\beta$, and $\overline{\alpha}\overline{\beta}$. Translating $\alpha \to \beta$ as $\alpha \wedge \overline{\mathsf{ab}_1} \supset \beta$ and circumscribing ab_1 over $(\alpha \wedge \overline{\mathsf{ab}_1} \supset \beta) \wedge (\overline{\alpha} \vee \overline{\beta})$ results in the theory $\overline{\alpha}$, which is satisfied in only two of the three MPD maximal models.

Figure 8 gives a correct, if somewhat less obvious, procedure for translating an MPD axiomatization into an equivalent circumscriptive theory. Essentially, the translation reifies the MPD defaults and minimizes the non-traversal of the accessibility relation. In particular, models not accessible from one another under MPD must also be circumscriptively incomparable.

Given the literals $\mathcal{L} = \{l_1, l_2, \ldots, l_n\}$, we generate a unique term $[\![[\mathcal{T}]]\!]$ for each truth assignment $\mathcal{T}$: $\circ_i \lambda_i$, where $\circ_i$ is alphabetically ordered string concatenation and $\lambda_i = l_i$ if $\mathcal{T} \models l_i$ and $\overline{l}_i$ otherwise. We then add axioms such that given the truth assignment $\mathcal{T}$, $\mathsf{ab}([\![[\mathcal{T}]]\!])$ is satisfied. For example, for $n = 3$ we add axioms like $(l_1 \wedge \overline{l_2} \wedge l_3) \supset \mathsf{ab}(l_1 \overline{l_2} l_3)$ for every distinct assignment to l_1, l_2, and l_3.

We capture the accessibility relation by adding axioms forcing the less preferable truth assignment (under MPD) to be abnormal for at least the same terms as the more preferable assignment. First, we rewrite the default rules to include specificity explicitly. Then, for each default of the form $\alpha \rightarrow \beta$, we add all the axioms corresponding to the following schema: $\mathsf{ab}([\![\ \alpha\ \wedge\ \beta\ \wedge\ \mathsf{X}]\!])\ \supset\ \mathsf{ab}([\![\ \alpha\ \wedge\ \beta\ \wedge \mathsf{X}]\!])$, where X varies over the 2^n possible truth assignments to the n literals not in α or β. Finally, we do predicate circumscription of $\mathsf{ab}(x)$ over the conjunction of the axioms we have added and the propositional theory (i.e. we find the models setwise minimal in the extent of ab). Since predicate circumscription derives models that are setwise minimal in the extent of the predicate being circumscribed, models not related by defaults will be incomparable.

As a result, the translation procedure in figure 8 yields a circumscriptive theory with minimal models equivalent to the maximal models of the original MPD theory. The initial axioms we add (step 1) make all the possible truth assignments incomparable (i.e. in each one, ab is true for a unique term). The axioms corresponding to defaults (step 3) define a set-inclusion relation corresponding exactly to the accessibility relation defined by the set of defaults. In other words, $\mathsf{ab}(\mathcal{T}) \supset \mathsf{ab}(\mathcal{T}')$ iff $\mathcal{T}'$ is accessible from $\mathcal{T}$. Thus $\mathcal{T}'$ is circumscriptively less than $\mathcal{T}$ just in case $\mathcal{T}'$ is MPD preferable to $\mathcal{T}$, and $\mathcal{T}'$ is setwise minimal in the extent of ab iff $\mathcal{T}'$ is a maximal model.

6.2 Model Preference Defaults and Default Logic

MPD is intended to capture some of the salient points of default reasoning. In [11], Selman and Kautz use the complexity of certain sublanguages of MPD to derive intractability results for Reiter's default logic (DL) [9]. We give a general translation procedure from MPD axiomatizations in $\mathcal{DH}_{\mathbf{a}}^{+}$ to DL. However, there are fundamental differences between MPD and DL that make general comparisons of MPD and DL axiomatizations difficult.

An MPD default rule $\alpha \rightarrow \beta$ corresponds roughly to the DL default $\frac{\alpha : \beta}{\beta}$. If there is a more specific (blocking) MPD default rule $\alpha \wedge \gamma \rightarrow \overline{\beta}$, we can make the specificity explicit, writing $\frac{\alpha \wedge \overline{\gamma} : \beta}{\beta}$ for $\alpha \rightarrow \beta$. But this comparison is loose at best: in general, there may be no way to translate an MPD axiomatization into DL so that the MPD maximal models are in a one-to-one correspondence with the DL extensions; nor does the reverse translation seem probable.

Since MPD always generates complete truth assignments, any MPD axiomatization will have at least one maximal model. In contrast, DL yields *extensions*, or partial world views, and there are some DL theories that have no extensions. For example, the DL extension of $D = \{\frac{\mathsf{a}:\mathsf{b}}{\mathsf{b}}\}, W = \{\}$[6] is empty and thus admits all four models for these two literals: ab, $\mathsf{a\overline{b}}$, $\mathsf{\overline{a}b}$, $\overline{\mathsf{ab}}$. The MPD axiomatization consisting of the single default rule $D = \{\mathsf{a} \rightarrow \mathsf{b}\}$, $T = \{\}$, limits the preferred models: $\mathsf{a\overline{b}}$ is not maximal, and MPD yields only ab, $\mathsf{\overline{a}b}$, $\overline{\mathsf{ab}}$.

In addition, while DL is in general nonmonotonic, the derivation process for any single extension is monotonic: Once β is in an extension, $\overline{\beta}$ can never be in *that* extension, although it may be in another extension of the same DL theory.[7] In contrast, the presence of a default $\rightarrow \beta$ in no way guarantees that β will be in any maximal model. For example, the theory $D = \{\rightarrow \mathsf{b},\ \mathsf{b} \rightarrow \mathsf{a},\ \mathsf{a} \rightarrow \mathsf{b}\}$ under specificity has a single maximal model, $\mathsf{a\overline{b}}$. MPD axiomatizations thus allow a default to "undercut" its own support.

For some theories, Selman and Kautz suggest that it may be possible to force DL to generate "complete" extensions through the use of additional *closure defaults*. These DL defaults apply even when the monotonic part of the DL theory is empty—i.e., when nothing is known for certain. For MPD axiomatizations in $\mathcal{DH}^+_{\mathsf{a}}$, since we can recursively enumerate the conditions under which a literal may be assigned a specific truth value, we can write the requisite closure defaults in a single pass through the literals. It is as yet uncertain whether this technique can be extended to the other acyclic sublanguages of MPD.

Consider the literal β. Let $\{\alpha_i \rightarrow \beta\}_i$ be all the MPD default rules with consequent β; let $\{\alpha_i \wedge \gamma_{ij} \rightarrow \overline{\beta}\}_{ij}$ be more specific default rules, respectively.

[6]A DL theory consists of a monotonic theory or "world" W, and a set of default rules D.

[7]Note that this is *not* the same as semi-monotonicity. Semi-monotonicity states that adding an additional default will only add to any of the original extensions; the monotonicity of derivation described here is for a fixed set of defaults, and simply states that a default cannot undercut its own support.

(Non-existent α_i, γ_{ij}, should be regarded as $\top$.) We add the closure rules:[8]

$$\frac{\bigwedge_i(\overline{\alpha_i} \vee (\bigwedge_j \gamma_{ij})) : \overline{\beta}}{\overline{\beta}}$$

This same procedure can be applied to $\overline{\beta}$, yielding defaults for β. This ensures that every literal will be assigned a truth value by forcing the assignment of that truth value whenever we can guarantee that no conflicting default will apply.

For example, consider the MPD axiomatization

$$T = \{\} \qquad\qquad D = \left\{ \begin{array}{c} \rightarrow \mathsf{a} \\ \mathsf{a} \rightarrow \mathsf{c} \\ \mathsf{a} \wedge \mathsf{b} \rightarrow \overline{\mathsf{c}} \end{array} \right\}$$

The translations of the defaults are

$$W = \{\} \qquad\qquad D = \left\{ \frac{:\mathsf{a}}{\mathsf{a}}, \ \frac{\mathsf{a}\wedge\overline{\mathsf{b}}:\mathsf{c}}{\mathsf{c}}, \ \frac{\mathsf{a}\wedge\mathsf{b}:\overline{\mathsf{c}}}{\overline{\mathsf{c}}} \right. \tag{10}$$

Specificity adds $\overline{\mathsf{b}}$ to the preconditions of the second default rule. In addition, we have the following closure defaults:

- For a, no defaults. ($\rightarrow \mathsf{a}$ overrides any default for $\overline{\mathsf{a}}$, and the first default is a closure for a.)

- For b,

$$\frac{:\mathsf{b}}{\mathsf{b}} \text{ and } \frac{:\overline{\mathsf{b}}}{\overline{\mathsf{b}}} \tag{11}$$

- For c, $\frac{\overline{\mathsf{a}}\vee\overline{\mathsf{b}}:\mathsf{c}}{\mathsf{c}}$ and $\frac{\overline{\mathsf{a}}\vee\mathsf{b}:\overline{\mathsf{c}}}{\overline{\mathsf{c}}}$, or

$$\left. \frac{\overline{\mathsf{a}} : \mathsf{c}}{\mathsf{c}}, \ \frac{\overline{\mathsf{b}} : \mathsf{c}}{\mathsf{c}}, \ \frac{\overline{\mathsf{a}} : \overline{\mathsf{c}}}{\overline{\mathsf{c}}}, \ \frac{\overline{\mathsf{b}} : \overline{\mathsf{c}}}{\overline{\mathsf{c}}} \right\} \tag{12}$$

The maximal models of the MPD axiomatization are $\mathsf{ab}\overline{\mathsf{c}}$ and $\mathsf{a}\overline{\mathsf{bc}}$; the extensions for the DL theory given by (10), (11), and (12) are the same.

[8] Disjuncts in the assumptions of DL defaults should be regarded as a syntactic sugar for multiple rules: $\frac{\beta\vee\gamma:\alpha}{\alpha}$ expands to $\frac{\beta:\alpha}{\alpha}$ and $\frac{\gamma:\alpha}{\alpha}$.

7 Summary

Selman and Kautz introduce the logic of model preference defaults (MPD) as a tool for the exploration of tradeoffs between expressivity and computational complexity in nonmonotonic reasoning. In [11, 12], they prove a number of complexity results for various restricted classes of MPD axiomatizations.

In this paper, we argue that the relevance of those results is restricted by the limitations of the formalism itself. MPD does not appear to have any straightforward and useful correspondence to two of the most important styles of nonmonotonic reasoning: default logic and circumscription. This makes it difficult to move from complexity results obtained for MPD to any more general conclusions.

In addition, we explore the applicability of MPD to three tractable problems in nonmonotonic reasoning: inheritance reasoning, temporal reasoning, and vivid knowledge bases. In each case, we show that the style of reasoning supported by MPD is at variance with the commonly accepted set of inferences. In all cases, this is due at least in part to the restriction of inference in MPD to the construction of a single complete truth assignment. This makes it impossible to represent incomplete knowledge or partial conclusions on the one hand, or to quantify over sets of conclusions (models or extensions) on the other.

Acknowledgments

We gratefully acknowledge assistance from Leora Morgenstern, Tom Dean, Dina Q. Goldin, Bart Selman, Henry Kautz, and various anonymous referees.

References

[1] David W. Etherington, Alex Borgida, Ronald J. Brachman, and Henry Kautz. Vivid knowledge and tractable reasoning. In *IJCAI–89* [3], pages 1146–1152.

[2] Steve Hanks and Drew V. McDermott. Nonmonotonic logic and temporal projection. *Artificial Intelligence*, 33:379–412, 1987.

[3] *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, Detroit, Michigan, August 1989. Morgan Kaufmann Publishers, Inc.

[4] Henry A. Kautz. The logic of persistence. In *Proceedings of the Fifth National Conference on Artificial Intelligence*, pages 401–405, Philadelphia, Pennsylvania, August 1986. Morgan Kaufmann Publishers, Inc.

[5] Henry A. Kautz and Bart Selman. Hard problems for simple default logics. *Artificial Intelligence*, 49(1–3):243–279, May 1991.

[6] Hector J. Levesque. Making believers out of computers. *Artificial Intelligence*, 30:81–108, 1986.

[7] John M. McCarthy. Circumscription—a form of nonmonotonic reasoning. *Artificial Intelligence*, 13(1,2):27–39, 1980.

[8] Drew V. McDermott and Jon Doyle. Non-monotonic logic I. *Artificial Intelligence*, 13:41–72, 1980.

[9] Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1,2):81–132, 1980.

[10] Raymond Reiter and Giovanni Criscuolo. On interacting defaults. In *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, pages 270–276, Vancouver, B.C., Canada, August 1981. Morgan Kaufmann Publishers, Inc.

[11] Bart Selman and Henry Kautz. The complexity of model-preference default theories. In *Proceedings of the Seventh Biennial Conference of the Canadian Society for the Computational Studies of Intelligence*, pages 102–109, Edmonton, Alberta, Canada, June 1988. Morgan Kaufmann Publishers, Inc.

[12] Bart Selman and Henry Kautz. Model-preference default theories. *Artificial Intelligence*, 45(3):287–322, 1990.

[13] Bart Selman and Hector Levesque. The tractability of path-based inheritance. In *IJCAI–89* [3], pages 1140–1145.

[14] Yoav Shoham. *Reasoning About Change: Time and Causation from the Standpoint of Artificial Intelligence.* MIT Press, Cambridge, Massachusetts, 1987.

[15] Lynn Andrea Stein. Resolving ambiguity in nonmonotonic inheritance hierarchies. *Artificial Intelligence*, 55(2–3):259–310, 1992.

[16] David S. Touretzky. *The Mathematics of Inheritance Systems.* Morgan Kaufmann Publishers, Inc., Los Altos, California, 1986.

[17] David S. Touretzky, John F. Horty, and Richmond H. Thomason. A clash of intuitions: The current state of nonmonotonic multiple inheritance systems. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, pages 476–482, Milan, Italy, August 1987. Morgan Kaufmann Publishers, Inc.