Investigations of Model-Preference Defaults

Mark Boddy
Robert P. Goldman
Keiji Kanazawa
Lynn Andrea Stein
Brown University Department of Computer Science
Providence, RI 02912, U.S.A.

Technical Report No. CS-89-13, March 1989

Investigations of Model-Preference Defaults

Mark Boddy
Robert P. Goldman
Keiji Kanazawa
Lynn Andrea Stein *
Department of Computer Science
Brown University
Providence, RI 02912

March 10, 1989

Abstract

In this paper, we explore the expressive and computational properties of the logic of *model-preference defaults* (MPD) [11, 12]. We extend the set of sublanguages of MPD known to be tractable. We compare MPD with the approaches to nonmonotonic reasoning represented by default logic [8] and circumscription [6]. Finally, we discuss MPD's suitability for various applications of nonmonotonic reasoning, including inheritance, temporal reasoning, and knowledge representation. In doing so, we show that MPD provides tractable solutions to problems of temporal projection and to mixing strict and defeasible inheritance.

1 Introduction

In [11], Selman and Kautz introduce *model-preference defaults* (MPD), a propositional default logic intended as a testbed for exploring nonmonotonic reasoning. They define the search problem of finding a maximal model given an MPD axiomatization. They show that for general MPD axiomatizations this search problem is NP-hard. They describe

*This work was supported by IBM Graduate Fellowships, National Science Foundation grants IST 8416034, IST 8515005 and IRI-8612644, Office of Naval Research grant N00014-79-C-0529, and by the Advanced Research Projects Agency of the Department of Defense, under Contract No. F49620-88-C-0132, monitored by the Air Force Office of Scientific Research.

a polynomial search algorithm for a sublanguage of MPD, limited to acyclic horn-form defaults, and with a propositional theory containing only literals.

These results raise some further questions: What are the expressive limitations of MPD? Can Selman and Kautz's results be used to show that there are other types of default reasoning with polynomial-time search problems? Defining maximal models in terms of defaults is intuitively appealing; do such models correspond to the models or extensions generated by other formalisms? In this paper, we address each of these questions in turn.

First, we explore the expressive and computational properties of MPD. We present new results concerning the tractability of the search problem for sublanguages of MPD. We show that finding a maximal model for any set of acyclic defaults with a propositional Horn theory can be accomplished in polynomial time. Unfortunately, Selman and Kautz's criterion for model maximality leads to anomalous results in the presence of monotonic Horn theories, including the counter-intuitive result that a seemingly stronger axiomatization has weaker consequences.

Next, we compare MPD with the approaches to nonmonotonic reasoning represented by default logic [8] and circumscription [6]. Differences in underlying assumptions appear to make it difficult, if not impossible, to translate between general MPD and default logic. On the other hand, there is a simple translation procedure for any MPD axiomatization into a circumscriptive theory which we prove yields the same set of models.

Finally, we discuss MPD's suitability for various applications of nonmonotonic reasoning, including inheritance, temporal reasoning, and knowledge representation. We show that MPD provides a polynomial algorithm for inheritance with strict and defeasible arcs, and for temporal projection using theories of the form defined by Shoham in his logic of chronological ignorance [13].

In the next section, we provide a precise definition for MPD and some of its sublanguages.

2 Background and Terminology

The description of MPD in this section does not parallel the definitions given in [11] and [12], but rather a corrected version [10]. The earlier definitions did not behave correctly in the presence of a non-empty monotonic theory. The difference lies in the careful distinction made between truth-assignments and models.

An MPD *axiomatization* is a pair $\langle T, D\rangle$. T is a set of formulas in the propositional calculus, referred to as the (monotonic or propositional) *theory*. D is a set of defaults of

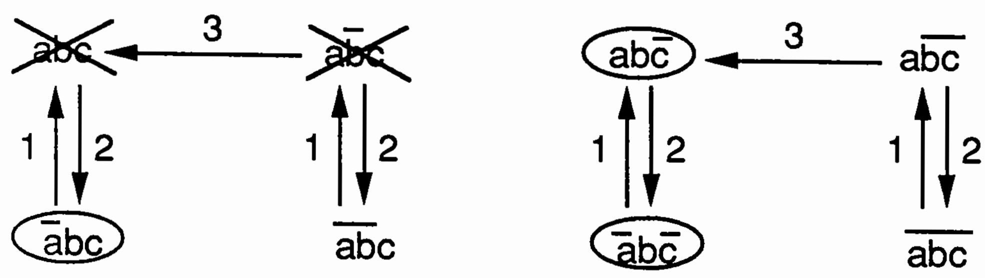

Figure 1: Preference relation graph

the form $\alpha \rightarrow \beta$, where α is a conjunction of literals, and β is a single literal. The defaults define a partial order, the preference relation, over the models of the propositional theory. Intuitively, MPD prefers models in which more default rules "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 $\mathcal{T}$. MPD axiomatizations may make use of a *specificity relation.*[1] 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 $D = \{1 :\rightarrow a,\ 2 :\rightarrow \overline{a},\ 3 : a \rightarrow b\}$, $T = \{\overline{a} \vee \overline{c}\}$. 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., $\overline{ab}c$ is not maximal, because $\overline{a}bc$ is both more preferable, and a model.

It will be useful to consider *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 a single, maximal model for a given axiomatization. $\mathcal{D}$ is MPD without specificity. $\mathcal{D}^+$ is $\mathcal{D}$ with the addition of the specificity relation. $\mathcal{D}_a^+$ restricts the defaults of $\mathcal{D}^+$ to be acyclic. $\mathcal{DH}$ is the set of MPD axiomatizations with "Horn-form" defaults: if $\alpha \rightarrow \beta$ is a

[1] A language with the specificity relation is marked with a + superscript, *e.g.*, $\mathcal{D}^+$.

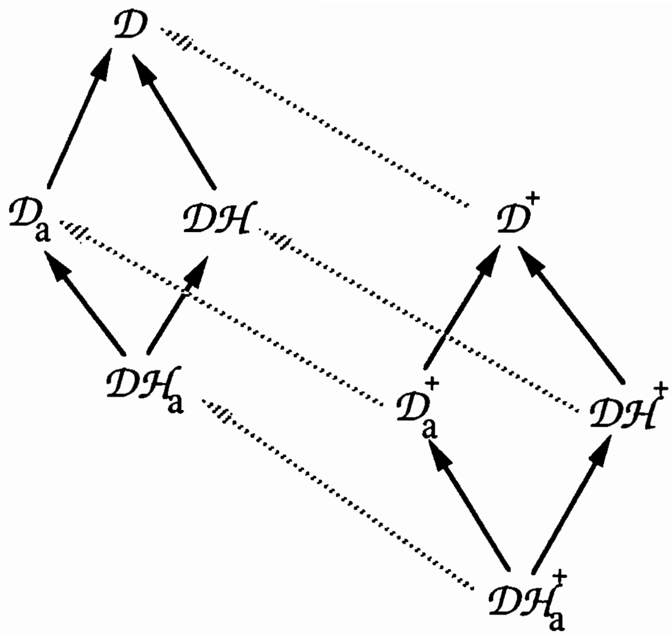

Figure 2: Language tree graph

default, then α contains no negated literals. $\mathcal{DH}^+$ adds specificity; $\mathcal{DH}_a$ and $\mathcal{DH}_a^+$ require that the defaults be acyclic and Horn-form. The tree of language containment relations is shown in Figure 2; dashed arrows represent the translation described in section 3.2

In [12], Selman and Kautz show that the search problems for $\mathcal{D}$, $\mathcal{D}^+$ and $\mathcal{DH}^+$ with empty theories are 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 NP-hard. Selman and Kautz describe a polynomial algorithm for $\mathcal{DH}_a^+$ with a monotonic theory made up of literals. In the next section, we show that it is possible to generalize the language $\mathcal{DH}_a^+$ in two ways without losing its favorable computational properties.

3 Complexity and Expressive Power

3.1 Tractable Sublanguages

Selman and Kautz show that the search problem for the language $\mathcal{DH}_a^+$ has a polynomial algorithm. $\mathcal{DH}_a^+$ requires Horn-form defaults, whose antecedents cannot contain negated literals. Selman and Kautz also hypothesize that there is a polynomial algorithm for the search problem for $\mathcal{DH}_a^+$ with monotonic Horn theories.

We define the languages $\mathcal{D}_a$ and $\mathcal{D}_a^+$, whose default rules must be acyclic, but need not be Horn-form. We provide a polynomial algorithm for $\mathcal{D}_a$ and $\mathcal{D}_a^+$ (which contains $\mathcal{DH}_a^+$) with monotonic Horn theories.

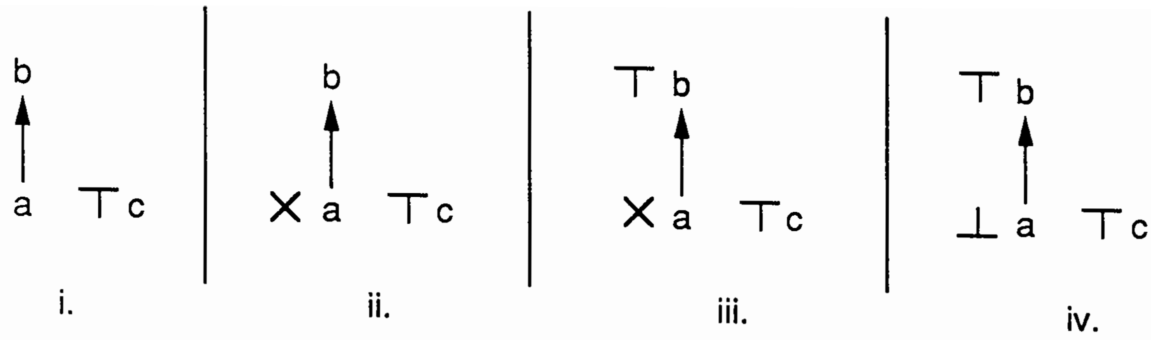

Figure 3: Literal dependency graph.

Max_model_$\mathcal{D}_a/\mathcal{D}_a^+$

1. Construct a directed graph with nodes corresponding to the set of literals. For each default rule, add arcs from each literal mentioned in the antecedent to the literal in the consequent.
2. Traverse the graph in topological order, updating the value for a literal l as follows:
   (a) Find all the applicable default rules. "Applicable" means l or $\bar{l}$ is the consequent, all the antecedents are satisfied (ambiguous ('x') satisfies either positive or negative antecedents), and the rule is not blocked.
   (b) If all applicable rules agree on a value, assign that value, if it is consistent with the propositional theory and the values already assigned to literals.
   (c) If there are conflicting rules, assign a value of 'x'.
   (d) If there are no applicable rules, assign an arbitrary value consistent with the propositional theory and the previously assigned values.
3. Make a second pass through the graph, replacing all the 'x' values with an arbitrary assignment consistent with the propositional theory and previous assignments.

A proof of correctness of this algorithm is given in appendix A.

Figure 3 illustrates the algorithm for the MPD axiomatization $D = \{1 :\rightarrow a,\ 2 :\rightarrow \bar{a},\ 3 : a \rightarrow b\}$, $T = \{\bar{a} \vee \bar{c}\}$. Figure 3 i. gives the literal graph after assigning $\top$ to c in step 2d. Since no defaults have consequent c (or $\bar{c}$), this asssignment is unconstrained. In Figure 3 ii., a has been assigned X, according to step 2c. Figure 3 iii. shows the assignment to b, which must occur after a has been marked. b is assigned $\top$ according to step 2b. Finally, Figure 3 iv. shows the complete assignment. Since c was assigned $\top$ in Figure 3 i., a must receive $\perp$ to be consistent with the monotonic theory, T. Had c been assigned $\perp$, a could now receive either $\top$ or $\perp$. The accessibility graph on truth assignments is shown in Figure 1.

3.2 Incorporating Specificity

It is possible to rewrite the defaults for a particular MPD axiomatization for one of the languages including specificity, so as to eliminate the need for specificity as a meta-linguistic property. Given the defaults $\rightarrow a$ and $b \rightarrow \overline{a}$, we rewrite the first as $\overline{b} \rightarrow a$. The rewritten set of defaults will generate the same model preferences as the original set, without requiring specificity.

More generally, given a set of defaults D, rewrite as follows: Divide D into equivalence classes, such that two defaults $d = \alpha \rightarrow \beta$ and $d' = \alpha' \rightarrow \beta'$ are members of the same class just in case β and β' contain the same literal (negated or not). Define a relation on each class, such that $d < d'$ just in case $\beta = \neg\beta'$ and $\alpha \subset \alpha'$. Traverse each class according to this relation, such that a rule d' is rewritten only if its successors have been rewritten. To rewrite a rule d: for each immediate successor d' and for each literal l s.t. $l \in \alpha' \wedge l \notin \alpha$, write a rule $\alpha \wedge \neg l \rightarrow \beta$ (invert the sense l had in α').

3.3 Aliasing

One of the motivations for extending $\mathcal{DH}_{\mathsf{a}}^{+}$ to $\mathcal{D}_{\mathsf{a}}^{+}$ 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—for example to represent a situation where, if we shoot at someone when the light is on, we typically kill them, while if we shoot at someone in the dark, we typically kill an innocent bystander. We might invent a dual predicate, dark, that is the opposite of light. We could then write the rules shoot $\wedge$ light $\rightarrow$ kill_target and shoot $\wedge$ dark $\rightarrow$ kill_bystander.

Unfortunately, there is no satisfactory way to relate the truth values of light and dark. Defining dark as logically equivalent to $\overline{\text{light}}$ 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 Horn-form propositional theories. Indeed, even using (intractable) $\mathcal{DH}^{+}$ defaults, such aliasing requires that dark *always* be used in place of $\overline{\text{light}}$, and light in place of $\overline{\text{dark}}$.

3.4 Propositional Theories

Selman and Kautz's criterion for model maximality leads to anomalous results in the presence of monotonic Horn theories. When we replace a default $d : \alpha \rightarrow \beta$ with the corresponding strict implication $\alpha \supset \beta$, an additional model may become maximal: the one in which neither α nor β is satisfied. This is because monotonic rules do not affect

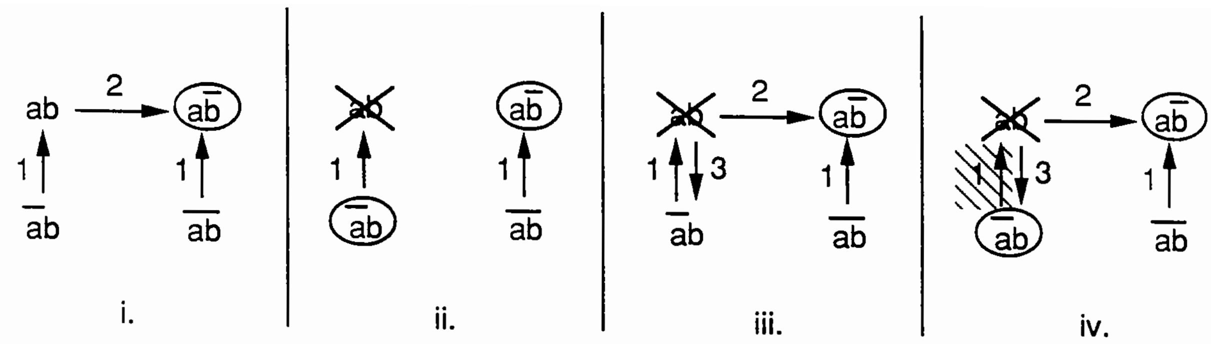

Figure 4: Accessibility graphs

model preferences and accessibility, so $d(T)$ is no longer accessible from T. For example, $D = \{1 :\rightarrow a,\ 2 : a \rightarrow \bar{b}\}$ $T = \{\}$ has maximal model $a\bar{b}$, while $D = \{1 :\rightarrow a\}$ $T = \{a \supset \bar{b}\}$ has *two* maximal models, $\bar{a}b$ and $a\bar{b}$, since the accessibility of $\{a\bar{b}\}$ from $\{ab\}$ by $2 : a \rightarrow b$ is removed. The accessibility graphs for these axiomatizations are shown in Figures 4-i and 4-ii, respectively.

We could try adding the default rule $\alpha \rightarrow \beta$ whenever $\alpha \supset \beta$ is in the monotonic theory. Unfortunately, there is in general no way to distinguish which of several defaults to add: given $a \supset \bar{b} = \bar{a} \vee \bar{b}$, we could add either $2 : a \rightarrow \bar{b}$ or $3 : b \rightarrow \bar{a}$. Adding all of the corresponding default rules may result in the same anomalous maximal models as not adding any of them. Figure 4-iii shows the result of adding both 2 and 3, and considering the result without specificity; Figure 4-iv shows the same axiomatization with specificity. The maximal models are circled, while truth assignments inconsistent with the propositional theory have been crossed out. We thus have the counter-intuitive result that a seemingly stronger axiomatization has weaker consequences. This problem 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 below)—the strategy of adding appropriate defaults may be more successful.

Another anomalous property of the accessibility relation in the presence of propositional theories is that, given a default $a \rightarrow b$, it is possible for a model where b is satisfied to be preferred to one where b is not, *even if there are no models of the theory that satisfy a.* Given a model in which the antecedents of a default are satisfied except for the value of a single literal l, the accessibility relation allows us to move from the original truth assignment, to a truth assignment where l is true and the consequent is false, to a truth assignment where l and the consequent are both true, back to a truth assignment where l is not true, but the consequent is. For example, given defaults $D = \{1 : j \rightarrow l,\ k \rightarrow \bar{l},\ l \rightarrow m\}$, $T = \{j, k, \bar{l}\}$, we get the (partial) preference relation

$$jk\bar{l}\bar{m} \rightleftharpoons jkl\bar{m} \rightarrow jklm \rightleftharpoons jk\bar{l}m$$

The second and third truth assignments are not models of the theory. Because of the definition of accessibility, the relation between them means that the fourth truth assignment is preferred to the first, with no particularly intuitive justification.

4 Comparisons to other formalisms

We investigate the relation between MPD and circumscription[6], and between MPD and Reiter's default logic[8]. We show that MPD can be translated directly into circumscription, but is fundamentally incompatible with default logic.

4.1 Circumscription

Circumscription results in models that are minimal in the extent of some predicate(s) (or more generally, in the satisfaction of some formula(e)). 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. We provide a translation, after showing that the most obvious way of translating defaults will 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 rule contraposes, while the MPD axiom does not. The default rule $\alpha \rightarrow \overline{\beta}$ 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 $\beta \rightarrow \overline{\alpha}$. 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 $\alpha \rightarrow \beta$ 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 $\alpha \rightarrow \beta$ 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 5 gives a procedure for translating an MPD axiomatization into an equivalent circumscriptive theory. The translation from MPD must preserve the properties 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 write a term for each truth assignment $\mathcal{T}$ as $\{\lambda_i\}$, where $\lambda_i = l_i$ 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,

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 $[\mathcal{T}] = \bigwedge_i \lambda_i$, and $[\![\mathcal{T}]\!] = \circ_i \lambda_i$, where $\circ$ is string concatenation and $\lambda_i = l_i$ if $\mathcal{T} \models l_i$ and $\bar{l}_i$ otherwise.
3. For each truth assignment $\mathcal{T}$, add the axiom $[\mathcal{T}] \supset \mathsf{ab}([\![\mathcal{T}]\!])$.
4. For each default of the form $\alpha \rightarrow \beta$, add the axioms corresponding to the following schema:

$$\mathsf{ab}(\alpha \cup \{\overline{\beta}\} \cup \mathsf{X}) \supset \mathsf{ab}(\alpha \cup \{\beta\} \cup \mathsf{X})$$

   where X varies over truth assignments to literals not in α, β.
5. Add as axioms the propositional theory.
6. Circumscribe $\mathsf{ab}(x)$ over the resulting theory.

Figure 5: Translation procedure

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, as discussed in section 3.2. Then, for each default of the form $\alpha \rightarrow \beta$, we add all the axioms corresponding to the following schema:

$$\mathsf{ab}(\alpha \cup \{\overline{\beta}\} \cup \mathsf{X}) \supset \mathsf{ab}(\alpha \cup \{\beta\} \cup \mathsf{X})$$

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.

A proof of correctness for this algorithm is given in appendix A.

4.2 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) [8]. We give a general translation procedure from

MPD axiomatizations in $\mathcal{DH}_a^+$ 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 $D = \{\frac{\alpha:\beta}{\beta}\}, W = \{\}$[2] is empty and thus admits all four models for these two literals: $\alpha\beta, \alpha\overline{\beta}, \overline{\alpha}\beta, \overline{\alpha\beta}$. The MPD axiomatization consisting of the single default rule $D = \{\alpha \rightarrow \beta\}, T = \{\}$, limits the preferred models: $\alpha\overline{\beta}$ is not maximal, and MPD yields only $\alpha\beta, \overline{\alpha}\beta, \overline{\alpha\beta}$.

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.[3] 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 $D = \{\rightarrow \beta,\ \beta \rightarrow \alpha, \alpha \rightarrow \beta\}$ under specificity has a single maximal model, $\alpha\overline{\beta}$. 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 $\mathcal{DH}_a^+$, 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 (non-existent α_i, γ_{ij}, should

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

[3] 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.

be regarded as $\top$). We add the closure rules:[4]

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

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

$$T = \{\} \qquad D = \begin{cases} \rightarrow \alpha \\ \alpha \rightarrow \gamma \\ \alpha \wedge \beta \rightarrow \overline{\gamma} \end{cases}$$

The translations of the defaults are

$$W = \{\} \qquad D = \left\{ \tfrac{:\alpha}{\alpha}, \tfrac{\alpha \wedge \overline{\beta} : \gamma}{\gamma}, \tfrac{\alpha \wedge \beta : \overline{\gamma}}{\overline{\gamma}} \right\} \tag{1}$$

Specificity adds $\overline{\beta}$ to the preconditions of the second default rule. In addition, we have the following closure defaults: for α, no defaults ($\rightarrow$ *alpha* overrides any default for $\overline{\alpha}$, and the first default is a closure for α); for β,

$$\tfrac{:\beta}{\beta} \text{ and } \tfrac{:\overline{\beta}}{\overline{\beta}} \tag{2}$$

for γ, $\frac{\overline{\alpha} \vee \overline{\beta} : \gamma}{\gamma}$ and $\frac{\overline{\alpha} \vee \beta : \overline{\gamma}}{\overline{\gamma}}$, or

$$\frac{\overline{\alpha} : \gamma}{\gamma}, \frac{\overline{\beta} : \gamma}{\gamma}, \frac{\overline{\alpha} : \overline{\gamma}}{\overline{\gamma}}, \frac{\overline{\beta} : \overline{\gamma}}{\overline{\gamma}} \Big\} \tag{3}$$

The maximal models of the MPD axiomatization are $\alpha\beta\overline{\gamma}$ and $\alpha\overline{\beta}\gamma$; the extensions for the DL theory given by (1), (2), and (3) are the same.

5 Applications

5.1 Inheritance

One motivation behind the design of MPD was to model defeasible inheritance. $\mathcal{DH}_{\mathbf{a}}^{+}$ generalizes *completely credulous inheritance.* This form of inheritance is on-path, credulous and upward-reasoning in the sense of [16]. It incorporates a pre-emption strategy [14] akin

[4] 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 $\frac{\gamma : \alpha}{\alpha}$.

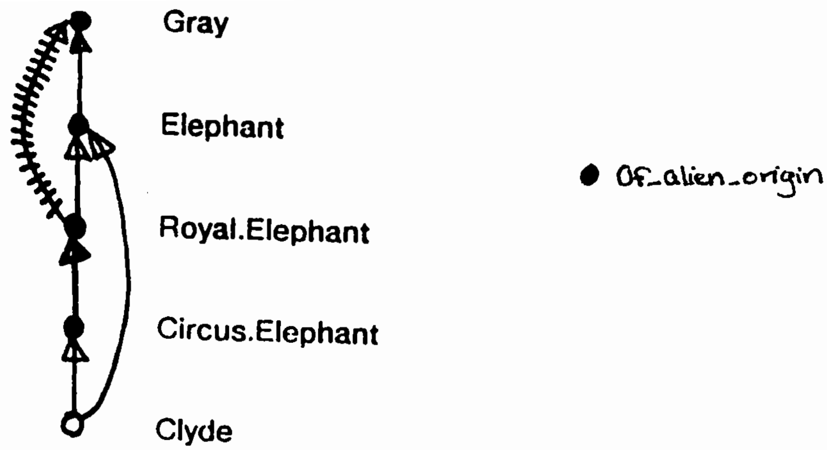

Figure 6: A simple inheritance network

to Touretzky's inferential distance [15]. This strategy of ambiguity resolution is captured by the specificity relation. Unlike credulous inheritance, *completely* credulous inheritance allows properties that are neither explicitly inherited nor explicitly excluded to be either inherited or not.

A non-monotonic inheritance network can be represented in MPD by writing defaults corresponding to the arcs of the network. To capture inferential distance, we resolve conflicting defaults according to the algorithm of Stein[14]. Figure 6 shows an example from Touretzky[15]. The corresponding defaults are:

$$clyde \rightarrow circus.elephant$$
$$circus.elephant \rightarrow royal.elephant$$
$$royal.elephant \rightarrow elephant$$
$$elephant \rightarrow gray$$
$$clyde \rightarrow elephant$$
$$royal.elephant, elephant \rightarrow \overline{gray}$$

To see the difference between this approach and that of standard credulous inheritance, consider the above network with an additional, disconnected, node: *of.alien.origin*. The MPD axiomatization of this network yields two maximal models: one in which Clyde is a space alien, and one in which he is not. A standard credulous inheritance reasoner would draw no conclusion as to whether Clyde is from outer space.

MPD (in particular, $\mathcal{DH}_a^+$ with horn theories) gives a polynomial algorithm for inheritance with strict and defeasible arcs. We represent a strict arc by a horn clause and a default rule. For example, a better representation of 'elephant world' would say that royal.elephants are by definition elephants. While in general it is not possible to determine which defaults to add to MPD axiomatizations with horn clauses to make the correct models maximal, it is possible to do so for an MPD axiomatization corresponding to an inheritance hierarchy: we translate x (strictly) ISA y as $x \supset y$ and $x \rightarrow y$.

5.2 Chronological Minimization

Shoham [13] 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 $\mathcal{D}_{\mathbf{a}}^{+}$.

Shoham begins with sentences of the form

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

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 ($\bigwedge_{i=1}^{n} \Diamond\overline{\psi_i}$), then we conclude the consequent ($\Box\omega$).

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

$$\bigwedge_{i=1}^{m} \varphi_i \rightarrow \omega \tag{4}$$

$$\text{For each } \Diamond\psi_j, \quad \bigwedge_{i=1}^{m} \varphi_i \wedge \psi_j \rightarrow \overline{\omega} \tag{5}$$

$$\text{For each } \Diamond\psi_j, \quad \rightarrow \overline{\psi_j} \tag{6}$$

Default rule (4) says that if the preconditions hold, we assume the consequent. (5) allows the qualifications to interfere: if any one of the qualifications holds, we can't assume the consequent.[5] In this case, (4) will be blocked by the more specific (5). Finally, (6) ensures that the qualifications do not, in general, hold. (6) 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 $\omega_{-1} = (t-1, p)$ and $\overline{\omega}_{-1} = (t-1, \overline{p})$. For each such state, we add a persistence rule:

$$\omega_{-1} \rightarrow \omega \tag{7}$$

[5]In fact, we assume the *negation* of the consequent. This is slightly stronger than what Shoham's rules yield, since the epistemic nature of his rules allows modal atoms whose truth values are indeterminate.

In general, *states persist*. We must now modify (4) to allow it to override this persistence. We do this by adding $\overline{\omega}_{-1}$ to the left-hand side:

$$\overline{\omega}_{-1} \wedge \bigwedge_{i=1}^{m} \varphi_i \rightarrow \omega \tag{8}$$

We make the same change to (5), allowing it to override (8):

$$\text{For each} \Diamond\psi_j, \quad \omega_{-1} \wedge \bigwedge_{i=1}^{m} \varphi_i \wedge \psi_j \rightarrow \overline{\omega} \tag{9}$$

(6) remains the same.

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

$$\rightarrow \overline{\omega} \tag{10}$$

Unlike states, actions and their negations are asymmetric. Unless we have reason to believe otherwise, *actions do not occur*, while it is difficult if not impossible to tell in general whether a state, or its negation, should "usually" be assumed to hold. Where it *is* possible, the qualifications of our causal rules dictate the sign of the state.

For example, 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 schemas:

$$(t, \mathsf{turn\_key}) \rightarrow (t+1, \mathsf{started\_motor}) \qquad \text{by (4)}$$
$$(t, \mathsf{turn\_key}) \wedge (t, \mathsf{potato\_in\_tailpipe}) \rightarrow \neg(t+1, \mathsf{started\_motor}) \qquad \text{by (5)}$$
$$\neg(t, \mathsf{potato\_in\_tailpipe}) \qquad \text{by (6)}$$

$$\neg(t, \mathsf{potato\_in\_tailpipe}) \wedge (t, \mathsf{put\_potato\_in\_tailpipe}) \rightarrow (t+1, \mathsf{potato\_in\_tailpipe}) \quad \text{by (8)}$$

$$(t, \mathsf{potato\_in\_tailpipe}_t) \rightarrow (t+1, \mathsf{potato\_in\_tailpipe})$$
$$\neg(t, \mathsf{potato\_in\_tailpipe}_t) \rightarrow \neg(t+1, \mathsf{potato\_in\_tailpipe})$$
$$(t, \mathsf{started\_motor}_t) \rightarrow (t+1, \mathsf{started\_motor})$$
$$\neg(t, \mathsf{started\_motor}_t) \rightarrow \neg(t+1, \mathsf{started\_motor}) \qquad \text{by (7)}$$

and

$$\neg(t, \mathsf{put\_potato\_in\_tailpipe})$$
$$\neg(t, \mathsf{turn\_key}) \qquad \text{by (10)}$$

5.3 The parking lot problem

In [4], 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.*, [6, 7, 8]). As Hanks and McDermott [3] 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, the most popular solution has been to incorporate the asymmetric structure of time into the process of selecting extensions to a given theory [4, 13]. 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 MPD, specifically in $\mathcal{D}^+$:

- $\mathsf{car}_t \rightarrow \mathsf{car}_{t+1}$, $\overline{\mathsf{car}_t} \rightarrow \overline{\mathsf{car}_{t+1}}$ – 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.
- $\overline{\mathsf{car}_{t+1}} \wedge \mathsf{car}_t \rightarrow \mathsf{steal}_t$ – explanation.

Suppose that our universe contains only the time points t_1, t_2, and t_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. Given the following theory, corresponding to observations made at the time points t_1 and t_3:

$$\{\mathsf{car}_1, \overline{\mathsf{car}_3}\}$$

There are two maximal models for this theory, corresponding to the car's getting stolen at t_1 or t_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).

5.4 Vivid reasoning

Selman and Kautz offer MPD as a logic suited for vivid reasoning. Levesque [5] defines vivid knowledge bases as databases representing all facts of interest in an explicit way.[6] 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 problem is to return a single, maximal model of an axiomatization, regardless of whether there is one such model, or many. This property of MPD requires making arbitrary choices, even when not choosing is the appropriate course of action, as in Reiter and Criscuolo's "Nixon diamond" [9]. Since only one model is computed, we are denied even the rough measure of certainty obtained by taking the union and intersection of credulous models. This results in a blurring of the distinctions not only between truth and likelihood, but among truth, likelihood, and arbitrary choice. Since MPD models do not indicate which conclusions are arbitrary, MPD is not suited for use in a hybrid reasoning system like that suggested by Etherington and Brachman. Such a program would be unable to tell when to refer a query to its more accurate, but more expensive, general problem-solver.

While it would be possible to rerun the algorithm, making random choices for undetermined literals, there may in general be exponentially many maximal models. In addition, to do so would be such a major redefinition of the MPD project as to be beyond the scope

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

of this paper. The MPD languages were designed to determine what could be done within tractable default languages. Further, an assumption underlying the design of the language is that there are problems which can be axiomatized such that arbitrary choices will not interfere with performance.

6 Conclusion

In this paper, we explore the expressive and computational properties of model-preference defaults. We show that the sublanguage of MPD for which Selman and Kautz have a polynomial algorithm can be extended to a more general language which allows the use of negation in default rules. It is also possible to add horn-form theories to MPD axiomatizations and retain tractability. Unfortunately, the use of horn-form theories may result in anomalous maximal models.

We investigate the relation between MPD and other formalisms for nonmonotonic reasoning. We show that MPD can be translated directly into circumscription, but is fundamentally incompatible with Reiter's default logic. Contrasting the straightforward translation from MPD to circumscription with the incompatibility of MPD and default logic underscores the differences between the assumptions of circumscriptive and fixpoint approaches to nonmonotonic reasoning[2].

We discuss MPD's suitability as a formalism for various kinds of default reasoning. 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 $\mathcal{D}_{\mathrm{a}}^{+}$, for which we have a polynomial search algorithm. MPDs with horn-form theories can be used to represent inheritance hierarchies with strict and defeasible links, yielding a polynomial time algorithm for inheritance. MPD axiomatizations do not distinguish between reasoned conclusions and arbitrary choices, making MPD best suited for applications where representing ambiguity is not required.

Acknowledgements

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

A Proofs of correctness

Max_model_$\mathcal{D}_{\mathrm{a}}/\mathcal{D}_{\mathrm{a}}^{+}$

At every node, all the literals that could possibly affect the assignment to l (*i.e.*, appear in the antecedent of relevant defaults) have been assigned values.

If no rule applies, there is a maximal model for either value of l (as long as that assignment is consistent with the theory). Truth-assignments consistent with the current (partial) truth-assignment which differ only in the value of l are incomparable, since there is no applicable default preferring either value.

If conflicting rules apply, then the current set of assignments can be extended by assigning l a value of either true or false. Since no further assignments to literals can affect the applicability of the conflicting rules, *any* two models consistent with the current truth assignment and differing only w.r.t. l will be mutually preferable.[7] So we assign an ambiguous value ('x'), and allow either negated or un-negated antecedents to be satisfied by 'x'; it is always possible to move to a truth assignment where the literal has the value required to apply a default, and then back, once the default has been applied. The propositional theory can only affect the final (second pass) assignment to l—the definition for the preference relation allows accessibility through truth assignments inconsistent with the propositional theory.

Because the monotonic theory is horn-form, whenever it is satisfiable, we can always construct a model incrementally, in any order we choose. In particular, we can follow the order dictated by this algorithm. For each literal l, we resolve together the propositional theory, previous assignments to literals, and a proposed value for l. If the conjunction is satisfiable, then there is a model for the theory consistent with the current truth assignment (including the value for l).

We construct a model in two passes, where unambiguous or arbitrary assignments made on the first pass can never be contradicted. The assignments to ambiguous values (made on the second pass) can be arbitrary (consistent with the propositional theory), since either model is always accessible from the other. Thus we are guaranteed to generate a model for the propositional theory. It is also guaranteed to be a maximal model, since for each literal we apply all possible defaults, and the acyclic nature of the defaults guarantee that later assignments to other literals will not affect the defaults that apply. □

MPD_to_Circumscription

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

[7] In fact, this has counterintuitive ramifications, which we discuss in section 3.4.

by the set of defaults. In other words, $\mathsf{ab}(M) \supset \mathsf{ab}(M')$ iff M' is accessible from M. Thus M' is circumscriptively less than M just in case M' is MPD preferable to M, and M' is setwise minimal in the extent of ab iff M' is a maximal model. □

References

[1] Ronald Brachman and David Etherington. Vivid knowledge and tractable reasoning. Unpublished abstract, 1988.

[2] David W. Etherington. Relating default logic and circumscription. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence*, pages 489–494, 1987.

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

[4] Henry Kautz. The logic of persistence. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pages 401–405, 1986.

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

[6] John McCarthy. Circumscription—a form of nonmonotonic reasoning. *Artificial Intelligence*, 13:27–39, 1980.

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

[8] Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

[9] Raymond Reiter and Giovanni Criscuolo. On interacting defaults. In *Proceedings of the Fourth International Joint Conference on Artificial Intelligence*, pages 270–276, 1981.

[10] Bart Selman, 1988. Personal Communication.

[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 Computational Studies of Intelligence*, pages 102–109, 1988.

[12] Bart Selman and Henry Kautz. Model preference default theories. Technical report, Department of Computer Science, University of Toronto, 1988.

[13] Yoav Shoham. *Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence.* PhD thesis, Yale University Department of Computer Science, 1987.

[14] Lynn Andrea Stein. Extensions as possible worlds. In *Formal Aspects of Semantic Networks.* Morgan Kaufmann.

[15] David Touretzky. *The Mathematics of Inheritance Systems.* Morgan Kaufmann Publishers, Inc., 1986.

[16] 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 10th International Joint Conference on Artificial Intelligence*, pages 476–482, 1987.