

Sometimes Updates Are Circumscription

Marianne Winslett
Computer Science Department
University of Illinois
Urbana, IL 61801

Abstract

Model-based revision of knowledge bases expressed as first-order theories was shown in [Winslett 88b] to be useful as a means of describing and reasoning about the effects of actions. This paper shows that model-based theory revision is actually expressible as a form of circumscription. This shows that in certain applications, the cumbersome conceptually machinery of circumscription can be replaced by the intuitively simpler ideas of model-based theory revision. Where theory revision techniques are insufficient to capture the semantics of change in an application, circumscription will offer a more flexible environment. In addition, future advances in computing circumscription can be mapped to improvements in computing theory revisions, and vice versa.

1 Introduction

There is a good deal of interest in the philosophical, AI, and database communities in the problem of revising logical theories (e.g., [Dalal 88, Fagin 83, 86, Forbus 89, Gärdenfors 88ab, Ginsberg 86, Oddie 78, Reiter 87, Satoh 88, Weber 86, Winslett 88a]). Section 2 of this paper reviews a model-based semantics for revising such theories [Winslett 88b]. This paper will not dwell on the uses for this semantics; rather, in Section 4 we will show that circumscription, a formal approach to common-sense reasoning, can be used to capture exactly this semantics of theory revision. We conclude with a discussion of the advantages and drawbacks of each approach.

2 Knowledge Base Updates

This section reviews the model-based method of theory revision presented in [Winslett 88b]. Let $\mathcal{T}$ be a theory describing some aspect of the world. Some of the formulas in $\mathcal{T}$ will describe data about the current state of the world, data that we will be willing to revise in the face of new information. Other formulas of $\mathcal{T}$, however, will describe beliefs that we are loathe to retract. For example, a formula giving the laws of thermodynamics or the rules of a game is a formula that we will likely be unwilling to retract no matter what new information arrives.

Such formulas are called *protected*, and the remaining formulas of $\mathcal{T}$ are *unprotected*.

In applications, even different unprotected beliefs are not held equally strongly. For example, given the information that a friend of mine is right now in Alaska, I may be more willing to change my beliefs about her vacation plans in order to explain that fact than to change my belief that she works in Silicon Valley. The predicates of $\mathcal{T}$ are assigned an a priori priority ordering, based on our willingness to retract beliefs in atoms over those predicates. More precisely, the predicates of $\mathcal{T}$ have been partitioned into nonempty strata numbered 1 through l , such that the priority of minimization of changes in predicates in stratum i is higher than for those in stratum j iff $i < j$.

Each model $\mathcal{M}$ of our theory $\mathcal{T}$ is a description of what we believe may be the current state of the world. Given new information about the world in the form of a formula α , we would like to update $\mathcal{M}$ by making α true in $\mathcal{M}$, while changing $\mathcal{M}$ as little as possible. Further, the protected formulas of $\mathcal{T}$ must still be true in the updated version of $\mathcal{M}$. As an additional wrinkle, the priority ordering on predicates means that changes in the truth valuations of atoms over some predicates are considered to be less important than changes in others, so that the measure of minimality of changes must consider exactly which predicates have been changed.

Intuitively, two models with the same universe agree on an formula if they assign it the same truth valuation; otherwise they differ on it. They agree on a predicate P if P has the same extension in both models.

Definition 1. Let α be a formula and $\mathcal{M}$ a model. $\text{Revise}(\mathcal{M}, \alpha)$ is the set of models $\mathcal{M}'$ such that

1. $\mathcal{M}'$ and $\mathcal{M}$ have the same universe and agree on all functions.
2. α and the protected formulas of $\mathcal{T}$ are true in $\mathcal{M}'$.
3. There is no other model $\mathcal{M}''$ such that for some $1 \leq i \leq l$,
 - $\mathcal{M}''$ satisfies (1) and (2);
 - $\mathcal{M}''$ and $\mathcal{M}'$ agree on all predicates in strata 1 through $i - 1$; and
 - the differences between $\mathcal{M}''$ and $\mathcal{M}$ on predicates in stratum i are a proper subset of the differences between $\mathcal{M}'$ and $\mathcal{M}$ on those predicates.

The models produced by incorporating a formula α into theory $\mathcal{T}$ are given by

$$\text{Revise}(\mathcal{T}, \alpha) = \bigcup_{\mathcal{M} \in \text{Models}(\mathcal{T})} \text{Revise}(\mathcal{M}, \alpha).$$

This definition does not say what formulas should appear in the revised theory; it simply tells what the models of that theory should be. A method of computing the result models has been devised and implemented by Forbus for use in reasoning about actions in the domain of qualitative physics [Forbus 89]. His theory revision approach incorporates an additional heuristic to choose between multiple possible result models: comparison of the cardinality of sets of proposed model changes. His implementation is described in [Forbus 89], included in these proceedings.

Example. As a running example, let us build a small description of a room, inspired by Ginsberg [87], and then revise it as we move around the objects in the room. The theory describes a living room containing a clock, a TV, and two air vents in the floor. The clock and TV can be either on the floor or on the vents; only one object fits on a vent at a time. The room is stuffy if both vents are blocked. This scenario is encoded by the formulas below. Protected formulas:

$$\begin{aligned} & \forall xyz. [\text{on}(x, y) \wedge \text{on}(x, z)] \rightarrow y = z \\ & \forall xyz. [\text{on}(x, y) \wedge \text{on}(z, y)] \rightarrow [x = z \vee y = \text{floor}] \\ & \forall xy. \text{on}(x, y) \rightarrow [(y = v1 \vee y = v2 \vee y = \text{floor}) \wedge \\ & \quad (x = \text{TV} \vee x = \text{clock})] \\ & \text{stuffy} \leftrightarrow \exists xy. \text{on}(x, v1) \wedge \text{on}(y, v2) \\ & \forall x \exists y. \text{on}(x, y) \vee (x = \text{floor}) \vee (x = v1) \vee (x = v2) \end{aligned}$$

These formulas ensure that (1) an object can only be in one place at a time, (2) only one object can be in a place at a time, (3) if x is on y then x is an object and y is a location, (4) the room is stuffy when both vents are blocked, and (5) every object is somewhere. The unprotected formulas of $\mathcal{T}$:

$$\begin{aligned} & \text{on}(\text{TV}, v1) \\ & \text{on}(\text{clock}, \text{floor}). \end{aligned}$$

“On” belongs in priority stratum 1, and “stuffy” in priority stratum 2. Intuitively, this means that whether the room is stuffy should be derived from the locations of objects, rather than the reverse. For simplicity, we will take a Herbrand Universe assumption for this example, so that the only objects are $v1$, $v2$, the TV, clock, and floor, and that all these constants are distinct. The only Herbrand model of $\mathcal{T}$, restricted to “on” and “stuffy”, is

$$\begin{aligned} & \text{on}(\text{TV}, v1) \\ & \text{on}(\text{clock}, \text{floor}). \end{aligned}$$

Suppose the clock is now moved from the floor to $v2$. If we update $\mathcal{T}$ by inserting $\text{on}(\text{clock}, v2)$, then the following result model is produced under Definition 1:

$$\begin{aligned} & \text{on}(\text{TV}, v1) \\ & \text{on}(\text{clock}, v2) \\ & \text{stuffy}. \end{aligned}$$

If “stuffy” were also in priority stratum 1, then a model where the TV moved to the floor and the room remained unstuffy would also be admitted. If “stuffy” were in stratum 1 and “on” were in stratum 2, then only this latter model would be produced.

3 Review of Prioritized Circumscription

The goal of circumscription [Etherington 88, McCarthy 87ab, Shoham 87] is to pick out the models of a theory $\mathcal{T}$ that are minimal with respect to a particular partial ordering. Models are ordered on the basis of the extents of certain of their predicates; in the ordering, models with smaller extents precede models with larger extents. We begin this section with the model-theoretic characterization of this intuition, and then present the second-order axiom that also captures the intuition.

Let S be the set of predicates whose extents are to be minimized. Assume that the predicates of S are partitioned into nonempty strata $S_1, \dots, S_l$, such that minimizing the extents of the predicates in S_i is of higher priority than minimizing those in S_j whenever $i < j$. A convenient notation for S is $S_1 < \dots < S_l$, as in $P_1, P_2 < P_3 < P_4, P_5$.

Recall that when comparing two models $\mathcal{M}_1$ and $\mathcal{M}_2$ on the extent of the predicates of stratum S_i , we may not care whether $\mathcal{M}_1$ and $\mathcal{M}_2$ agree on a set V_i of other predicates, in which case the predicates of V_i are said to be “allowed to vary.” If $S = S_1 < \dots < S_l$ gives the predicates to be minimized, then we will assume that when minimizing at stratum S_i , predicates of S_j are allowed to vary, for all $j > i$. Additional predicates not in S may also be allowed to vary when minimizing at S_i ; $V = V_1 < \dots < V_l$ is a convenient notation for recording the identity of these additional predicates. Under this notation, every predicate in V_j , for $j \geq i$, is allowed to vary when minimizing the predicates of S_i . The equality predicate cannot appear in either S or V .

Definition 2 presents the partial order on models that is introduced by this variant of circumscription. (Throughout this paper, $\bar{x}$ will be a tuple of terms, of the appropriate arity for the function or predicate taking $\bar{x}$ as an argument.)

Definition 2. A model $\mathcal{M}_1$ is preferred to model $\mathcal{M}_2$ at the i th stratum (written $\mathcal{M}_1 <_i \mathcal{M}_2$) iff

1. $\mathcal{M}_1$ and $\mathcal{M}_2$ have identical universes;
2. $\mathcal{M}_1$ and $\mathcal{M}_2$ agree on all predicates and functions, except possibly those of S_i and V_i ;
3. For all predicates P in stratum S_i , and for all $\bar{x}$ such that $P(\bar{x})$ is true in $\mathcal{M}_1$, $P(\bar{x})$ is true in $\mathcal{M}_2$; and
4. For some predicate P in stratum S_i and some $\bar{x}$, $P(\bar{x})$ is false in $\mathcal{M}_1$ and true in $\mathcal{M}_2$.

Intuitively, circumscription picks out those models that are minimal with respect to $<$. $\mathcal{M}_1$ is preferred to $\mathcal{M}_2$ (written $\mathcal{M}_1 < \mathcal{M}_2$) iff

1. For some $1 \leq i \leq l$, $\mathcal{M}_1 <_i \mathcal{M}_2$; and
2. For all $1 \leq j < i$, $\mathcal{M}_2 \not<_j \mathcal{M}_1$.

Finally, a model $\mathcal{M}$ of $\mathcal{T}$ is preferred if no model of $\mathcal{T}$ is preferred to $\mathcal{M}$.

The second-order characterization of circumscription says, intuitively speaking, that the predicates of S cannot be replaced by other predicates such that $\mathcal{T}$ is still true and the resulting models are “smaller”. For a theory $\mathcal{T}$, $\text{Circ}(\mathcal{T}; S; V)$ is defined to be the theory

$$\mathcal{T}(S; V) \wedge \forall S' \forall V' \neg [(S' < S) \wedge \mathcal{T}(S'; V')]. \quad (1)$$

In formula (1), $\mathcal{T}(S; V)$ is just $\mathcal{T}$, and $\mathcal{T}(S'; V')$ is the theory obtained by replacing all occurrences of the predicates of S and V by their counterparts in S' and V' , throughout $\mathcal{T}$. The second-order construct S' is obtained from S by replacing each predicate P that appears in S by a predicate variable P' with the same arity as P . V' is constructed from V in the same manner.

In formula (1), the formula $S' < S$ is shorthand for

$$\exists i. (1 \leq i \leq l) \wedge \forall j. (1 \leq j \leq i) \rightarrow \left[\bigwedge_{P \in S_j} \forall \bar{x} (P'(\bar{x}) \rightarrow P(\bar{x})) \wedge \bigvee_{P \in S_i} \exists \bar{x} (P(\bar{x}) \wedge \neg P'(\bar{x})) \right], \quad (2)$$

a proof-theoretic counterpart of our usual minimality criterion.

4 Relation of Updates to Circumscription

This section shows that the model-based semantics of Section 2 for updating logical theories can be expressed using prioritized circumscription.

In recent work, Satoh has [Satoh 88] independently proposed a model-based semantics for theory revision that is similar to ours¹; interestingly, Satoh uses second-order logic to define his semantics. Though he does not present his semantics as a type of circumscription, the form of his definition makes the close ties of his update semantics to circumscription readily apparent.

This paper only considers the case of a single update; the general case, where n updates are applied in succession to $\mathcal{T}$, uses a different, more complex type of circumscription, and is described in detail in [Winslett 88c]. Intuitively, the goal of the circumscription is to minimize the differences between each model $\mathcal{M}$ of $\mathcal{T}$ and the models in $\text{Revise}(\mathcal{M}, \alpha)$. The transformation of updates into circumscription is realized by first slightly rewriting the theory to be revised, then defining within the rewritten theory the update to be performed, and finally circumscribing the changes wrought by that update. These three phases are described below, followed by a proof of correctness for the transformations.

Before rewriting $\mathcal{T}$, we must change the language $\mathcal{L}$ underlying $\mathcal{T}$ to contain a different set of predicates. The new language $\mathcal{L}'$ contains three predicates in place of every non-equality predicate P of the old language: *oldP*, *newP*, and *changedP*, all with the same arity as P . Intuitively, this new language will be used to talk about what was true before and after the update, and also about what changed when the update occurred.

We next create the theory $\text{rewrite}(\mathcal{T}, \alpha)$ by applying the following four steps to $\mathcal{T}$.

1. Replace all occurrences of non-equality predicates P by *oldP*, throughout $\mathcal{T}$.

¹The distinction is that Satoh judges the minimality of a set of changes in a model $\mathcal{M}$ of $\mathcal{T}$ by a global comparison with all other sets of changes in models of $\mathcal{T}$ (sharing the same universe and function interpretations) that make the inserted formula true. We judge the minimality of a set of changes by comparison only with other sets of changes in $\mathcal{M}$.

2. Add additional protected formulas: make a second copy of the protected formulas, replacing each occurrence of a non-equality predicate *oldP* by *newP*. Intuitively, the new protected formulas require that both the old and new states of the world satisfy the protected formulas.
3. Suppose that the update to be performed is the insertion of a formula α . Represent this information by the unprotected formula $\text{new}(\alpha)$, which is α with each occurrence of a non-equality predicate P replaced by *newP*.
4. The new predicates *changedP* are to be true exactly when the truth of an atom of a non-equality predicate P alters in a model when α is inserted. In other words, add the new protected formulas

$$\forall \bar{x}. \text{changedP}(\bar{x}) \leftrightarrow [\text{oldP}(\bar{x}) \neq \text{newP}(\bar{x})]. \quad (3)$$

The logical connective “ $\neq$ ” is exclusive or, and is not related to the equality predicate.

$\text{Rewrite}(\mathcal{T}, \alpha)$ is the theory to be circumscribed. The circumscription policy is to circumscribe the *changedP* predicates in update stratum order. When minimizing the predicates of a particular stratum, the *changedP* and *newP* predicates of equal or higher strata are irrelevant, and therefore should be allowed to vary. More precisely, if the predicates of $\mathcal{T}$ are partitioned into l nonempty strata for update purposes, then $S = S_1 < \dots < S_l$, where $S_i = \{\text{changedP} \mid P \text{ is in stratum } i\}$. Similarly, $V = V_1 < \dots < V_l$, where $V_i = \{\text{newP} \mid P \text{ is in stratum } i\}$. The circumscription to be performed is given by $\text{Circ}(\text{rewrite}(\mathcal{T}, \alpha); S; V)$.

Example, continued. $\text{Rewrite}(\mathcal{T}, \text{on}(\text{clock}, v2))$ is the theory

$$\begin{aligned} & \forall xyz. [\text{oldOn}(x, y) \wedge \text{oldOn}(x, z)] \rightarrow y = z \\ & \forall xyz. [\text{oldOn}(x, y) \wedge \text{oldOn}(z, y)] \rightarrow [x = z \vee y = \text{floor}] \\ & \forall xy. \text{oldOn}(x, y) \rightarrow [(y = v1 \vee y = v2 \vee y = \text{floor}) \wedge (x = \text{TV} \vee x = \text{clock})] \\ & \text{oldStuffy} \leftrightarrow \exists xy. \text{oldOn}(x, v1) \wedge \text{oldOn}(y, v2) \\ & \forall x \exists y. \text{oldOn}(x, y) \vee (x = \text{floor}) \vee (x = v1) \vee (x = v2) \\ & \forall xyz. [\text{newOn}(x, y) \wedge \text{newOn}(x, z)] \rightarrow y = z \\ & \forall xyz. [\text{newOn}(x, y) \wedge \text{newOn}(z, y)] \rightarrow [x = z \vee y = \text{floor}] \\ & \forall xy. \text{newOn}(x, y) \rightarrow [(y = v1 \vee y = v2 \vee y = \text{floor}) \wedge (x = \text{TV} \vee x = \text{clock})] \\ & \text{newStuffy} \leftrightarrow \exists xy. \text{newOn}(x, v1) \wedge \text{newOn}(y, v2) \\ & \forall x \exists y. \text{newOn}(x, y) \vee (x = \text{floor}) \vee (x = v1) \vee (x = v2) \\ & \text{oldOn}(\text{TV}, v1) \\ & \text{oldOn}(\text{clock}, \text{floor}) \\ & \text{newOn}(\text{clock}, v2) \\ & \forall x \forall y. \text{changedOn}(x, y) \leftrightarrow [\text{oldOn}(x, y) \neq \text{newOn}(x, y)] \\ & \text{changedStuffy} \leftrightarrow [\text{oldStuffy} \neq \text{newStuffy}] \end{aligned}$$

In this example,

$$S = \text{changedOn} < \text{changedStuffy},$$

and

$$V = \text{newOn} < \text{newStuffy}.$$

The two Herbrand models of $\text{rewrite}(\mathcal{T}, \alpha)$ are

Model 1	Model 2
oldOn(TV, v1)	oldOn(TV, v1)
oldOn(clock, floor)	oldOn(clock, floor)
newOn(clock, v2)	newOn(clock, v2)
changedOn(clock, floor)	changedOn(clock, floor)
changedOn(clock, v2)	changedOn(clock, v2)
newOn(TV, v1)	newOn(TV, floor)
newStuffy	changedOn(TV, v1)
changedStuffy	changedOn(TV, floor).

According to the semantics for theory revision, the result of moving the TV to v1 should be model 1. Using the model-theoretic definition of circumscription, model 1 is preferred to model 2 at stratum 1 because `changedOn(TV, v1)` is true in model 2 and false in model 1, and there is no atom of “`changedOn`” that is true in model 1 and false in model 2. Therefore model 1 is the preferred model.

Theorem 1 describes the relationship between $\text{Circ}(\text{rewrite}(\mathcal{T}, \alpha); S; V)$ and $\text{Revise}(\mathcal{T}, \alpha)$:

Theorem 1. *Let $\mathcal{T}$ be a theory, and let α and ϕ be formulas. Then*

$$\begin{aligned} \text{Revise}(\mathcal{T}, \alpha) \models \phi \\ \text{iff} \\ \text{Circ}(\text{rewrite}(\mathcal{T}, \alpha); S; V) \models \text{new}(\phi). \end{aligned}$$

Proof. We sketch the proof in both directions. For the forward implication, let $\mathcal{M}$ be a model of $\mathcal{T}$, and let $\mathcal{M}'$ be a model in $\text{Revise}(\mathcal{M}, \alpha)$. We will now construct a model of $\text{CR}(\mathcal{T})$ (shorthand for $\text{Circ}(\text{rewrite}(\mathcal{T}); S; V)$) that corresponds to the pair $\mathcal{M}, \mathcal{M}'$.

Let $\mathcal{N}$ be a structure with the same universe and function interpretations as $\mathcal{M}$ and $\mathcal{M}'$, but over the language of $\text{CR}(\mathcal{T})$. For each nonequality predicate P , let $\text{oldP}(\bar{x})$ be true in $\mathcal{N}$ iff $P(\bar{x})$ is true in $\mathcal{M}$. For each nonequality predicate P , let $\text{newP}(\bar{x})$ be true in $\mathcal{N}$ iff $P(\bar{x})$ is true in $\mathcal{M}'$. Let $\text{changedP}(\bar{x})$ be true in $\mathcal{N}$ exactly when the truth valuations of $\text{oldP}(\bar{x})$ and $\text{newP}(\bar{x})$ differ in $\mathcal{N}$. Then for any formula ϕ , by construction ϕ is true in $\mathcal{M}'$ iff $\text{new}(\phi)$ is true in $\mathcal{N}$. Further, $\mathcal{N}$ satisfies all the formulas of $\text{rewrite}(\mathcal{T})$.

It remains to show that $\mathcal{N}$ is a preferred model of $\text{rewrite}(\mathcal{T})$. Suppose not. Then there is a model $\mathcal{M}'$ of $\text{rewrite}(\mathcal{T})$ that is preferred to $\mathcal{N}$. $\mathcal{N}$ and $\mathcal{M}'$ must agree on all “old” predicates, and there must exist a priority level i such that $\mathcal{N}$ and $\mathcal{M}'$ agree on all predicates of strata below i . In addition, the cases where $\text{changedP}(\bar{x})$ is true in $\mathcal{M}'$, over all choices of P at priority stratum i , must be a proper subset of the cases in which changedP is true in $\mathcal{N}$ at that stratum. But then one can construct a model $\mathcal{O}$ with the same universe and function interpretations as $\mathcal{M}$, and where $P(\bar{x})$ is true in $\mathcal{O}$ iff $\text{newP}(\bar{x})$ is true in $\mathcal{M}'$. The relationship between $\mathcal{M}$ and $\mathcal{O}$ shows that $\mathcal{M}'$ must not have been a member of $\text{Revise}(\mathcal{M}, \alpha)$ after all, as the differences between $\mathcal{O}$ and $\mathcal{M}$ at priority stratum i are a proper subset of the differences between $\mathcal{M}'$ and $\mathcal{M}$ at that stratum, and the differences at lower strata are identical. As this contradicts the assumption that $\mathcal{M}'$ was a member of $\text{Revise}(\mathcal{M}, \alpha)$, we conclude that the theorem holds in the forward direction.

For the reverse direction, suppose that $\mathcal{N}$ is a model of $\text{CR}(\mathcal{T})$, and $\text{new}(\phi)$ is true in $\mathcal{N}$. Let $\mathcal{M}$ be a model

with the same universe and function interpretations as $\mathcal{N}$ but over the language of $\mathcal{T}$, and in which $P(\bar{x})$ is true iff $\text{oldP}(\bar{x})$ is true in $\mathcal{N}$, for all nonequality predicates P . Let $\mathcal{M}'$ be a model with the same universe and function interpretations as $\mathcal{M}$, differing from $\mathcal{M}$ only in that if $\text{changedP}(\bar{x})$ is true in $\mathcal{N}$, then $\mathcal{M}$ and $\mathcal{M}'$ differ on the truth valuation of $P(\bar{x})$. By construction, $\mathcal{M}$ satisfies all the formulas of $\mathcal{T}$. If $\mathcal{M}'$ is not a member of $\text{Revise}(\mathcal{M}, \alpha)$, then there must be a model $\mathcal{O}$ in $\text{Revise}(\mathcal{M}, \alpha)$ such that $\mathcal{O}$ and $\mathcal{M}'$ agree on all predicates below some priority stratum i , and at priority stratum i the the points at which $\mathcal{O}$ and $\mathcal{M}$ differ on truth valuation are a proper subset of the points at which $\mathcal{M}'$ and $\mathcal{M}$ differ. But then one can construct a model $\mathcal{M}'$ of $\text{CR}(\mathcal{T})$, corresponding to the pair $\mathcal{M}, \mathcal{O}$, such that $\mathcal{M}'$ is preferred to $\mathcal{N}$. We conclude that in fact $\mathcal{M}'$ is in $\text{Revise}(\mathcal{M}, \alpha)$. Finally, since $\text{new}(\phi)$ is true in $\mathcal{N}$, by construction ϕ is true in $\mathcal{M}'$. $\diamond$

A variant of Theorem 1 was shown in [Winslett 88c] for the case where a series of updates is applied to $\mathcal{T}$.

5 Comparison of the Two Approaches

Given that model-based theory revisions can be expressed in circumscription, how should one choose between the two formalisms? In general, circumscription is more expressive, but also more complex.

Operator semantics. The semantics of operators in a reasoning-about-action application can be directly represented in a circumscriptive theory; under the theory revision approach, operator semantics is external to the theory.

Minimization policy. Priorities for minimization of different predicates can be directly represented in a circumscriptive theory; under the theory revision approach, any priorities were specified external to the theory.

Complexity. I feel that model-based theory revision is much simpler than circumscription for a human to understand, especially when one moves to a variant of circumscription that is sufficiently powerful to represent a series of updates.

Algorithms. The theory revision approach should enjoy computational advantages over circumscription, because the theory revision approach does not retain any information about previous states of the world. Circumscription based on situation calculus remembers all facts about every state of the world encountered. A hybrid architecture (retaining only those bits of information about prior situations that have an impact on the current situation) should perform better than either approach alone; see [Forbus 89] for a description of such an implementation.

6 Conclusions

We have shown how a model-based semantics for theory revision can be captured exactly using circumscription. This paper has shown how to translate a single update into circumscription; a series of updates can also be expressed using an appropriate type of circumscription [Winslett 88c].

This result captures an interesting connection between work on updates and circumscription. Now that the con-

nection has been established, future work on methods of computing circumscription and updates can be shared between the two disciplines. In addition, for a particular application, researchers can choose between the two paradigms based on considerations of expressiveness (circumscription is more expressive) and simplicity (model-based theory updates are conceptually simpler).

7 Acknowledgments

I would like to thank Moshe Vardi, Hirofumi Katsuno, and Peter Rathmann for their helpful comments. Vladimir Lifschitz suggested the simplified presentation of these results that I have used here.

8 References

- [Dalal 88] M. Dalal, "Investigations into a Theory of Knowledge Base Revision: Preliminary Report," AAAI-88, Minneapolis, August 1988.
- [Etherington 88] D. W. Etherington, *Reasoning With Incomplete Information*, Morgan Kaufmann, Los Altos, 1988.
- [Fagin 83] R. Fagin, J. D. Ullman, and M. Y. Vardi, "On the Semantics of Updates in Databases," *ACM Symposium on Principles of Database Systems*, April 1983.
- [Fagin 86] R. Fagin, G. M. Kuper, J. D. Ullman, and M. Y. Vardi, "Updating Logical Databases," in *Advances in Computing Research*, **3**, 1986.
- [Forbus 89] K. D. Forbus, "Introducing Actions into Qualitative Simulation," *Proceedings of the International Joint Conference on Artificial Intelligence*, Detroit, August 1989.
- [Gärdenfors 88a] P. Gärdenfors and D. Makinson, "Revisions of Knowledge Systems Using Epistemic Entrenchment," *Proc. Conf. on Theoretical Aspects of Reasoning About Knowledge*, Asilomar, March 1988.
- [Gärdenfors 88b] P. Gärdenfors, *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, Bradford Books, MIT Press, 1988.
- [Ginsberg 86] M. L. Ginsberg, "Counterfactuals," *Art. Intelligence*, **30**:1, 1986.
- [Ginsberg 87] M. L. Ginsberg and D. E. Smith, "Reasoning About Action I: A Possible Worlds Approach," in *Readings in Nonmonotonic Reasoning*, M. L. Ginsberg, ed., Morgan Kaufmann, Los Altos, 1987.
- [McCarthy 87a] J. McCarthy, "Circumscription—A Form of Non-Monotonic Reasoning," in *Readings in Nonmonotonic Reasoning*, M. L. Ginsberg, ed., Morgan Kaufmann, Los Altos, 1987.
- [McCarthy 87b] J. McCarthy, "Applications of Circumscription to Formalizing Common-Sense Knowledge," in *Readings in Nonmonotonic Reasoning*, M. L. Ginsberg, ed., Morgan Kaufmann, Los Altos, 1987.
- [Oddie 78] G. Oddie, "Verisimilitude," from *The Logic and Philosophy of Scientific Change*, in *Acta Philosophica Fennica*, **30**:2-4, 1978.
- [Pollack 76] N. Pollack, *Subjunctive Reasoning*, Reidel, Dordrecht, 1976.
- [Reiter 87] R. Reiter, "A Theory of Diagnosis from First Principles," *Artificial Intelligence*, April 1987.
- [Satoh 88] K. Satoh, "Nonmonotonic Reasoning by Minimal Belief Revision," *Proceedings of the International Conference on Fifth Generation Computer Systems 1988*, ICOT, December 1988.
- [Shoham 87] Y. Shoham, "A Semantical Approach to Nonmonotonic Logics," in *Readings in Nonmonotonic Reasoning*, M. L. Ginsberg, ed., Morgan Kaufmann, Los Altos, 1987.
- [Weber 86] A. Weber, "Updating Propositional Formulas," *Proc. of the Expert Database Systems Conference*, Charleston SC, April 1986.
- [Winslett 88a] M. Winslett, "A Framework for Comparison of Update Semantics," *ACM Symposium on Principles of Database Systems*, March 1988.
- [Winslett 88b] M. Winslett, "Reasoning About Action with a Possible Models Approach," AAAI-88, Minneapolis, August 1988. A full version appears as TR UIUCDCS-R-88-1428, Department of Computer Science, Univ. of Illinois—Urbana, May 1988.
- [Winslett 88c] M. Winslett, "Setwise Circumscription," unpublished manuscript, 27 pp., May 1988.