Semantics of First Order Belief Revision Based on Circumscription *

Li Yan Yuan and Jia-Huai You
Department of Computing Science
University of Alberta
Edmonton, Canada T6G 2H1
e-mail: yuan@cs.ualberta.ca, you@cs.ualberta.ca

Abstract

The *AGM postulates* for knowledge base revision are a set of rationality postulates that revision operations on knowledge systems should satisfy. Defining an appropriate semantics for belief revision that satisfies all of the postulates has been considered a challenging problem.

In this paper we present a novel application of circumscription to defining a semantics of belief revision systems. First, a first-order knowledge base is represented by a set of formulas in a first order epistemic belief language that contains objective propositions as well as belief propositions. Secondly, we define a revision semantics by applying a form of priority circumscription to the belief representation of the knowledge base. We prove that the semantics defined in this way satisfies the AGM postulates that are reformulated in our belief language.

Key Words: - knowledge base revision
- AGM postulates
- nonmonotonic reasoning, circumscription

*This paper has been printed as a Technical Report TR 92-17, Department of Computing Science, Univerisity of Alberta.

1 Introduction

A knowledge base should be updated as our perception of the world described by it changes. Revision is the most common type of updates: it adds newly acquired knowledge to the system. If the new knowledge is consistent with current beliefs, the revision is simple – just add it to the system. However, if the new knowledge is inconsistent with current beliefs, the conflict must be resolved somehow, usually by derogating some old beliefs. The question is how to choose a subset of the old beliefs as victims, and decide on a set of criteria for doing so.

Alchourron, Gardenfors, and Makinson have proposed a set of postulates, called the AGM postulates, which are based on well justified philosophical ground and provide a foundation for knowledge base revision [1]. Most previous proposals treat belief revision as a change operation over a set of propositional sentences (see, for example, [12, 7]), where both the representation and the semantics of knowledge systems are defined by the same set of logic sentences. It has been noticed that this approach, though provides a unified point of view on knowledge systems, lacks retrospective power, the main reason for its not being able to satisfy the AGM postulates.

Alchourron et al. proposed the *partial meet* revision semantics which satisfies the AGM postulates. The idea is to take the intersection of all possible candidate theories that reflect minimal changes and that can be resulted from conflict resolving. This semantics is not considered very realistic since by taking the intersection of all such theories, useful information may be lost and in many cases the resulting theory is simply the empty one. A more realistic semantics has been suggested by Fagin, Ullman, and Vardi [3] (henceforth the FUV semantics), which takes the disjunction of all such candidate theories. As indicated in [7], the FUV semantics is syntax dependent and therefore fails to satisfy the AGM postulates.

The revision problem over a theory in its syntactic form other than its closure is discussed by Nebel [12], where it is called *base revision*. Nebel discovered that base revision with epistemic relevance does not satisfy all of the AGM postulates and identified the cases where they are satisfied.

Belief revision in its essence is a meta level concept and operation. Therefore, it is natural to use some type of meta language to describe change operations. Recently, semantics for belief revision have been formulated in terms of some type of modal systems [2, 4]. In these approaches, reasoning about changes in ones beliefs reduces to model checking of certain meta level sentences. Both approaches in [2, 4] have been proved to satisfy all of the AGM postulates.

In this paper, we propose a new approach to the semantics of belief revision, in which a knowledge base is represented by a set of formulas in a belief language that contains objective as well as belief propositions, and the semantics is defined by applying a form of circumscription to the representation of the knowledge base. We show that all of the AGM postulates, if reformulated in our framework, are satisfied.

There are two basic premises on which our approach is based. First, like Nebel [12], revision in our approach is carried out over theories in their syntactic form other than their closure, i.e., deductively closed set. We argue that the notion of deductively closed set does not provide an appropriate framework for knowledge base revision.

Secondly, we will consider knowledge bases with an *arbitrary* binary relation, augmented by transitivity, over the set of all sentences in the underlying language. Such a binary relation will be called a *priority relation* in this paper; this is due to the common realization by the researchers in the field that knowledge revision must respect the epistemic importance in a knowledge base. As shown in [5, 12, 3], the AGM postulates can be satisfied by some revision semantics if the priority relation representing the epistemic importance satisfies certain conditions. These conditions actually express special cases of partial ordering. Thus the assumption of priority relation makes our approach more general.

Although our approach also relies on the concept of meta level objects, which we have called belief propositions, it is quite different from the work in [2, 4] in that our formulation of revision semantics is not semantically based on or dependent upon any modal logic. As a matter of fact, our approach can be formulated entirely in a conventional first order language. This is because a belief proposition $\mathcal{L}\phi$ in our epistemic belief language is treated as a *named* object, not as applying some modal operator $\mathcal{L}$ to ϕ, and thus can be viewed as, or simply replaced by, a *distinct* objective proposition. Thus, the problem of belief revision in our approach reduces to the better-known problem of reasoning with minimal models. An advantage of this is that revision operations can be realized directly on top of a circumscription algorithm (see, for example, [13]).[1]

That our approach does not depend on any modal logic, plus the fact that our approach is based on circumscription of first order theories, permits us to define our revision semantics for first order theories rather than propositional theories. This is another significant difference with all the other revision semantics in the literature.

The paper is organized as follows. In the next section we will carefully define and explain the belief language used in this paper; we will use an example to illustrate this language as well as the main idea in our approach. Since the original definition of priority circumscription is defined only for complete pre-ordering, in Section 3 we present an extension of priority circumscription that can perform minimization according to a partial ordering. Section 4 introduces the AGM postulates. In Section 5 we introduce a new framework for knowledge base revision and reformulate the AGM postulates to suit our new framework. Then in Section 6 we define a revision semantics and show it satisfies the reformulated AGM postulates.

[1] However, this is by no means to imply that circumscription is an easy problem.

2 The Belief Language

In this paper we assume a first order *belief* language L, which is a usual first order language that contains a set of objective predicate symbols and a set of belief predicate symbols. A usual first order formula is referred to as an *objective formula.* In the language of autoepistemic logic [11], belief formulas are of the form $\mathcal{L}\phi$ where ϕ is a formula and $\mathcal{L}$ is a special symbol in the alphabet of the language. For the purpose of this paper, we only need belief predicates. The name of an n-nary belief predicate consists of two parts: a normal predicate symbol and a prefix $\mathcal{L}$, such as $\mathcal{L}p(x_1, ..., x_n)$, where $\mathcal{L}p$ stands for a single predicate name. A belief predicate name, such as $\mathcal{L}p$, can technically be replaced by a *distinct* (or *reserved*) predicate symbol, and as such, whether the language contains the special symbol $\mathcal{L}$ or not is technically insignificant. The language can be equally defined as a purely first order language with some distinct predicate symbols. For this reason, in the rest of paper we will not use the usual belief symbol $\mathcal{L}$ but Greek letters such as α, β, and γ to denote belief predicates or belief propositions (i.e., 0-ary belief predicates). We will consider in this paper first order theories that consist of sentences with universally quantified variables. We often omit these quantifiers in sentences with the understanding that all free variables therein are universally quantified.

To illustrate the main idea in our approach, let us consider the following example.

Example 2.1 Consider a knowledge base expressed by the following set of formulas

$$K = \{bird,\ fly \leftarrow bird\}.$$

Should we later observe $\neg fly$, we need to revise our knowledge base K by the newly acquired knowledge $\neg fly$. Simply adding $\neg fly$ to K would result in an inconsistent set K':

$$K' = \{bird,\ fly \leftarrow bird,\ \neg fly\}.$$

To resolve inconsistency, one can remove a minimal amount of sentences so that the remaining sentences are consistent. These type of subsets have been called *maximum consistent subsets* in the literature. For K' above we obtain two maximum consistent subsets that contain $\neg fly$:

$$\Delta_1 = \{bird,\ \neg fly\}$$
$$\Delta_2 = \{fly \leftarrow bird,\ \neg fly\}.$$

By the FVU method, the revision semantics is then defined by taking the disjunction of all such candidate theories.

In our approach, each sentence ϕ in a knowledge base is represented by a formula $\phi \leftarrow \alpha(x_1, ..., x_n)$ where x_i occurs in ϕ and α is a distinct belief predicate intuitively meaning ϕ is believed. As usual, all the variables in the formula are universally quantified. Further, if no variables occur in ϕ then α is simply a belief proposition. The revised system of the above is then represented by the following set of sentences:

$$T_{K'} = \{bird \leftarrow \alpha,\ (fly \leftarrow bird) \leftarrow \beta,\ \neg fly \leftarrow \gamma\}.$$

Note that although K' is inconsistent its belief representation $T_{K'}$ is consistent.

Then the semantics of the system can be defined by applying circumscription to *maximize* the belief propositions (i.e., minimize the negations of the belief propositions) with γ having higher priority to be maximized, which yields

$$T_{K'} \wedge \gamma \wedge (\alpha \vee \beta) \wedge (\neg\alpha \vee \neg\beta).$$

Note that this expression implies $\neg fly$ and either $bird$ or $(fly \leftarrow bird)$ but not both at the same time. □

For the above example, the semantics defined in terms of maximizing beliefs is essentially the same as that of taking the disjunction of all candidate theories [3], as far as logical consequences of objective formulas are concerned. However, since the extension of belief propositions can be maximized according to a given priority and the circumscription *does not eliminate* any sentences from the belief representation of a theory, our revision semantics possesses retrospective power and therefore, can satisfy the AGM postulates.

3 Priority Circumscription Based on Partial Ordering

McCarthy introduced *circumscription* to express the idea that the extension of abnormal predicates should be minimized [9, 10]. Let $A(P, Z)$ be a first order theory, where P and Z are disjoint sets of predicates in A, and M and N be two models of A. Then we say N is (P, Z)-smaller than M if both models have the same extension over all predicates other than P and Z, but the extension of the predicates from P in N is a proper subset of that in M; and we say N is (P, Z)-minimal if no model of A is (P, Z)-smaller than N. Then $CIR(A; P; Z)$, the circumscription of A on P with variable Z, denotes a second order formula whose models are all (P, Z)-minimal models of A. Furthermore, priority circumscription $CIR(A; P^1 \succ \cdots \succ P^n; Z)$, where P^i's are partitions of P, is used to represent the idea that the extension of predicates from P^1 should have higher priority to be minimized than that of P^2 and the extension of predicates from P^2 have higher priority than that of P^3, etc.

As we mentioned earlier, our approach is based on maximizing belief propositions. The mechanism is called *maximizing circumscription*. This notion can be precisely defined.

Let $A(P)$ be a theory in a belief language, where P is the set of belief predicates whose extension is to be maximized. Then the *maximizing circumscription* of A on P, denoted as $MCIR(A;P)$, is defined as

$$MCIR(A;P) \equiv A(P) \wedge \neg\exists\mathcal{P}(A(\mathcal{P}) \wedge (\mathcal{P} > P))$$

where $\mathcal{P} > P$ means the extension of predicates from $\mathcal{P}$ is a proper superset of that from P. Maximizing circumscription can also be formalized in terms of circumscription on the negations of those predicates. For convenience, we may use $CIR(A;\neg P)$ to denote $MCIR(A;P)$, and $MCIR(A;\neg P)$ to denote $CIR(A;P)$. The priority version of maximizing circumscription $MCIR(A;P_1 \succ \cdots \succ P_n;Z)$ is similarly defined.

Lifschitz has shown that priority circumscription can be represented by parallel circumscription [8], that is, given a first order theory A and disjoint sets $P^1,...,P^n,Z$ of predicate symbols,

$$CIR(A;P^1 \succ ... \succ P^n;Z) = \bigwedge_{i=1}^{n} CIR(A;P^i;P^{i+1},...,P^n,Z).$$

The priority relation in the priority circumscription above is a linear, total relation amongst all predicate blocks of P^i's. In real applications, however, many priority relations are partial ordering, not total ordering, and priority circumscription cannot directly be used to express such minimization based on partial ordering. An extension of priority circumscription into partial ordering is given below.

Let P be a set of predicates, and $\preceq$ be a binary relation amongst P, augmented by obvious transitive closure. The binary relation is used to represent the priority relation amongst P, that is, $a \preceq b$ implies that b has at least as high priority as a to be minimized, and when $a \preceq b$ and $b \not\preceq a$ then b is considered having higher priority than a to be minimized. A partition $\{P^1,...,P^n\}$ of P is $\preceq$-compatible if it is defined by the equivalence relation that $a \equiv b$ if and only if $a \preceq b$ and $b \preceq a$, that is, $\{a,b\} \subseteq P^i$, for some i, if and only if $a \preceq b$ and $b \preceq a$. Obviously, for any given relation $\preceq$, its $\preceq$-compatible partition is unique. Furthermore, for each P^i, we define $LOW(P^i)$ as the set of all predicates in P that have lower priority than anyone in P^i to be minimized according to $\preceq$; that is, $LOW(P^i) = \{p \mid p \prec a \text{ for some } a \in P^i\}$. Now, we define

Definition 3.1 Let $A(P,Z)$ be a theory, where P and Z are disjoint sets of predicates in A, $\preceq$ be a priority relation defined over P, and $\{P^1,...,P^n\}$ be a $\preceq$-compatible partition. Then the $\preceq$-based priority circumscription is defined as

$$CIR(A(P,Z);P|\preceq;Z) = \bigwedge_{i=1}^{n} CIR(A(P,Z);P^i;(LOW(P^i)\cup Z)).$$

$MCIR(A(P,Z);P|\preceq;Z)$ is defined similarly. $\Box$

Priority circumscription $CIR(A;P^1 \succ ... \succ P^n;Z)$ is just a special case of $\preceq$-based circumscription when $\preceq$ is a linear order on $\{P^1,...,P^n\}$.

4 The AGM Postulates

In the framework of AGM, revision is an operation over deductively closed sets in the language of propositional logic [1].

Given a theory Γ, the deductively closed set of Γ is defined as the closure $\{\phi \mid \Gamma \vdash \phi\}$.

Let K be a deductively closed set, μ and ν be consistent sentences. The revision of K by μ, denoted as $K\hat{+}\mu$, represents a new knowledge system obtained from K by adding new knowledge represented in μ. Then the AGM postulates for revision are as follows:

(P1) $K\hat{+}\mu$ is a deductively closed set;

(P2) $K\hat{+}\mu \models \mu$;

(P3) $K \wedge \mu \models K\hat{+}\mu$;

(P4) $K\hat{+}\mu \models K \wedge \mu$ if $K \wedge \mu$ is consistent;

(P5) $K\hat{+}\mu$ is consistent;

(P6) $K\hat{+}\mu \equiv K\hat{+}\nu$ if $\mu \equiv \nu$;

(P7) $(K\hat{+}\mu) \wedge \nu \models K\hat{+}(\mu \wedge \nu)$;

(P8) $K\hat{+}(\mu \wedge \nu) \models (K\hat{+}\mu) \wedge \nu$ if $(K\hat{+}\mu) \wedge \nu$ is consistent.

The first postulate states that the revision of a deductively closed set must result in a deductively closed set. The second states that the new knowledge must be retained in the revision. P3 implies that the revision must be contained in the range of the simple union of old and new knowledge. The fourth represents the idea that the revision is done by simply adding μ to K if K is consistent with μ. P5 requires that the revision be consistent. The sixth specifies the principle of irrelevance of syntax.

The seventh, similar to P3, states that the revision of K by $\mu \wedge \nu$ must be subsumed by $K\hat{+}\mu$ augmented by ν. The last one, together with P7, states that if $K\hat{+}\mu$ is consistent with ν, then $(K\hat{+}\mu) \wedge \nu$ is equivalent to $K\hat{+}(\mu \wedge \nu)$.

There is a significant drawback in using the notion of deductively closed sets in the context of belief revision. Consider the following example, from [14], of revising the database

$$K = \{a, b\}$$

with $\neg a$. Intuitively, the new knowledge $\neg a$ should overwrite the old knowledge a, and this results in the revised knowledge base as $\{\neg a, b\}$. However, the deductively closed set of the database is

$$\{a, b, a \vee b, a \wedge b, a \leftarrow b, b \leftarrow a, ...\},$$

of which there are two maximum subsets that are consistent with $\neg a$, that is,

$$\{b, a \vee b, b \leftarrow a, ...\} \text{ and } \{a \leftarrow b, b \leftarrow a, ...\}.$$

By adding $\neg a$ into the two sets respectively, the first one is the deductively closed set of $\{b, \neg a\}$ and the second is that of $\{\neg b, \neg a\}$.

In order to accommodate $\neg a$ in the above knowledge base, derogating a is necessary. However, by adopting the framework of deductively closed set, we throw out both a and b. The reason that undesirable results are produced is that when the closure of a theory is calculated, anything that is a logical consequence is treated as important as those in the original theory. Humans are not logically omniscient, and so it is unrealistic to entail logic omniscience in any formal belief system [6].

Therefore, in this paper we use the notion of *knowledge set* instead of deductively closed set. A *knowledge set* is a finite set of first order sentences which may or may not be consistent, together with a priority relation (i.e., an arbitrary binary relation, closed under transitivity) over all sentences of the underlying language.[2]

We now use an example to show that even in the framework of knowledge sets we just described, the FUV method [3] and the like, which are based on taking the disjunction of all maximum consistent theories, fail to satisfy the AGM postulates.

Example 4.1 Consider the following knowledge set:

$$K = \{a,\ a \leftarrow b \wedge c,\ b,\ c,\ c \leftarrow d,\ d\}.$$

Assume the following priority relation $\preceq$:

$$\begin{array}{l} (b) \preceq (a \leftarrow b \wedge c) \\ (c) \preceq (a \leftarrow b \wedge c) \\ (d) \preceq (c \leftarrow d) \preceq (a \leftarrow b \wedge c) \end{array}$$

[2]Note that the question of what priority relations are useful, or even meaningful, is not the focus of this paper (but see [5]).

and $\mu = \{\neg a\}$ and $\nu = \{\neg d\}$ are new sentences to be added.

Suppose we revise K by μ. Since $K \cup \{\neg a\}$ is inconsistent, we need to obtain maximum consistent subsets of $K \cup \{\neg a\}$; i.e., the consistent subsets of $K \cup \{\neg a\}$ with a minimal amount of sentences removed and with less important sentence(s) removed first. To illustrate, since $\{\neg a\}$ is the most recent knowledge and should be included in any maximum consistent subset. Now $(a \leftarrow b \wedge c)$ has the next highest priority to be retained. Then, retaining both b and c would result in inconsistency, and thus, because b and c are not related by $\preceq$, either b or c but not both can be retained. This gives two ways of removing minimal amount of sentences. Should b be retained, either d or $(c \leftarrow d)$ should be removed to maintain consistency. This results in d being removed since $(c \leftarrow d)$ has the higher priority to be retained. We thus get two maximum consistent subsets under the priority relation given above:

$$\Gamma_1 = \{\neg a,\ a \leftarrow b \wedge c,\ b,\ c \leftarrow d\}$$

$$\Gamma_2 = \{\neg a,\ a \leftarrow b \wedge c,\ c,\ c \leftarrow d,\ d\}$$

By the FUV method, we get the disjunction of the above two subsets, i.e.,

$$K \hat{+} \mu = \{\neg a,\ a \leftarrow b \wedge c,\ c \leftarrow d,\ b \vee c,\ b \vee d\}.$$

Thus

$$(K \hat{+} \mu) \wedge \nu = \{\neg a,\ \neg d,\ a \leftarrow b \wedge c,\ c \leftarrow d,\ b\}.$$

On the other hand, by a similar process, K revised by $(\mu \wedge \nu)$ also has two maximum consistent subsets:

$$\Delta_1 = \{\neg a,\ \neg d,\ a \leftarrow b \wedge c,\ b,\ c \leftarrow d\}$$

$$\Delta_2 = \{\neg a,\ \neg d,\ a \leftarrow b \wedge c,\ c,\ c \leftarrow d\}.$$

Taking the disjunction, we get

$$K \hat{+} (\mu \wedge \nu) = \{\neg a,\ \neg d,\ a \leftarrow b \wedge c,\ c \leftarrow d,\ b \vee c\}.$$

Now we have $K \hat{+} (\mu \wedge \nu) \not\models (K \hat{+} \mu) \wedge \nu$, violating postulate P8. Therefore, the FUV method and the like, which are based on taking disjunctive theory of maximum consistent subsets, do not satisfy the AGM postulates. □

5 Representation and Semantics of Knowledge Systems

In this section, we show a new framework to represent knowledge systems and define their semantics. We then reformulate the AGM postulates to suit our new framework.

Let $K(P)$ be a knowledge set with priority relation $\preceq$, where P is the set of all objective predicates in language L. Then K can be represented as follows: For each sentence t in K, a belief predicate $\pi_t(x_1, ..., x_n)$ is introduced, where π_t is a distinct *new* predicate symbol not in P, and $x_1, ..., x_n$ are universally quantified variables occurring in t. Let B be the set of all such belief predicates. Then $T_K(P, B)$, called the *belief theory* of $K(P)$, is defined as

$$T_K(P, B) = \{t \leftarrow \pi_t(x_1, ..., x_n) \mid t \in K \text{ and } x_i \text{ occurs in } t, 1 \leq i \leq n\},$$

with the priority relation $\preceq$ carried over to B, i.e., $\alpha_1 \preceq \alpha_2$ iff t_1 and t_2 are in K such that $t_1 \preceq t_2$, and $t_1 \leftarrow \alpha_1$ and $t_2 \leftarrow \alpha_2$ are in T_K. For simplicity, we will denote $T_K(P, B)$ simply by $\{t \leftarrow \pi_t \mid t \in K\}$. Thus, π_t may represent, depending on the context, a belief predicate symbol in B, or a belief predicate with variables in a sentence.

Example 5.1 Consider

$$K = \{a,\ b \leftarrow a,\ \neg b\}$$

with the priority relation

$$\preceq = \{\{a\} \preceq \{b \leftarrow a\} \preceq \{\neg b\}\}.$$

Then

$$T_K(P, B) = \{a \leftarrow \alpha,\ b \leftarrow a \wedge \beta,\ \neg b \leftarrow \gamma\},$$

where $B = \{\alpha, \beta, \gamma\}$ and $\preceq = \{\alpha \preceq \beta \preceq \gamma\}$. Note that $b \leftarrow a \wedge \beta \equiv (b \leftarrow a) \leftarrow \beta$. $\Box$

Example 5.2 Let

$$K = \{bird(penguin),\ fly(x) \leftarrow bird(x),\ \neg fly(x) \leftarrow x = penguin\}$$

be a knowledge set with the following priority relation:

$$\begin{array}{l} bird(penguin) \preceq (\neg fly(x) \leftarrow x = penguin) \\ (fly(x) \leftarrow bird(x)) \preceq (\neg fly(x) \leftarrow x = penguin) \end{array}$$

Then $T_K(P, B)$ contains

$$\begin{array}{l} bird(penguin) \leftarrow \alpha \\ (fly(x) \leftarrow bird(x)) \leftarrow \beta(x) \\ (\neg fly(x) \leftarrow x = penguin) \leftarrow \gamma(x) \end{array}$$

where $B = \{\alpha, \beta(x), \gamma(x)\}$ and $\preceq = \{\alpha \preceq \gamma(x),\ \beta(x) \preceq \gamma(x)\}$. $\Box$

Definition 5.1 Let K be a knowledge set with priority relation $\preceq$, and $T_K(P, B)$ be the corresponding belief theory. Then the semantics of K, denoted as $F_{sem}(K, \preceq)$, or $F_{sem}(K)$ if $\preceq$ is understood, is defined by

$$MCIR(T_K(P, B); B| \preceq; P).$$

An objective formula ϕ is true in a knowledge system K if and only if $F_{sem}(K)$ logically implies ϕ. $\square$

Example 5.3 Consider, for example, the K and $T_K(P, B)$ in Example 5.1. We then have

$$F_{sem}(K) = MCIR(T(P, B); B| \preceq; P) \equiv T_K \wedge \gamma \wedge \beta,$$

which implies $\neg b$ and $(a \leftarrow b)$. Note that to maintain consistency, either a or $(a \leftarrow b)$ may be removed. However, according to the priority relation a should be removed since it has lower priority to survive.

For the K and $T_K(P, B)$ in Example 5.2, we have

$$F_{sem}(K) = MCIR(T(P, B); B| \preceq; P) \equiv$$

$$T_K \wedge \forall x\gamma(x) \wedge \forall x(x \neq penguin \rightarrow \beta(x)) \wedge (\alpha \vee \beta(penguin)) \wedge (\neg\alpha \vee \neg\beta(penguin)).$$

The formula $\forall x\gamma(x)$ holds because the belief predicate γ is maximized with the highest priority. The belief predicates β and α are unrelated in the priority relation. Conflict arises only when both $\beta(penguin)$ and α attempt to hold true. That is, the maximal extension of $\beta(x)$ includes any $\beta(t)$ where t is not $penguin$. This results in $\forall x(x \neq penguin \rightarrow \beta(x))$. When x is $penguin$, either $\beta(penguin)$ or α, but not both, holds true. This is expressed by $(\alpha \vee \beta(penguin)) \wedge (\neg\alpha \vee \neg\beta(penguin))$. $\square$

We thus have established a framework for knowledge revision, a knowledge system is represented by its belief theory, and the semantics of the system is determined by its "belief semantics" of maximizing the belief predicates. Therefore, all query evaluations toward the knowledge system should be directed to the belief semantics. However, there are two possible ways to view revision operations.

Suppose K is the given knowledge base with a priority relation $\preceq$. Consider the revision requests $\phi_1, ..., \phi_n$ in that order. With each revision request, the priority relation is enhenced so that the most current one always has the highest priority to survive. For the above revision sequence, let us denote the corresponding priority relations as $\preceq_1, ..., \preceq_n$.

In the first view, F_{sem} is taken purely as an operator, repeatedly applying to a knowledge set. In this view, a knowledge system's revolution with the above revision requests, given K and $\preceq$, can be described as the following sequence of the knowledge systems:

$$\begin{aligned} K_0 &= F_{sem}(K, \preceq) \\ K_1 &= F_{sem}(K_0 \cup \{\phi_1\}, \preceq_1) \\ K_2 &= F_{sem}(K_1 \cup \{\phi_2\}, \preceq_2) \\ &\ldots\ldots \\ K_n &= F_{sem}(K_{n-1} \cup \{\phi_n\}, \preceq_n) \end{aligned}$$

Note that in this view, the knowledge system and its semantics are uniformly presented to the user of the system.

In the second view, F_{sem} is treated as a mapping from an underlying physical system to a semantically meaningful knowledge system on which user's queries are evaluated against. More precisely, for each revision request ϕ_i, we simply add ϕ_i to the previous knowledge set. The underlying system contains a set of sentences, which may or may not be consistent; it is the mapping F_{sem} that interprets the system and provides the semantics. This can be described as the sequence:

$$\begin{aligned} K_0 &= F_{sem}(K, \preceq) \\ K_1 &= F_{sem}(K \cup \{\phi_1\}, \preceq_1) \\ K_2 &= F_{sem}(K \cup \{\phi_1, \phi_2\}, \preceq_2) \\ &\ldots\ldots \\ K_n &= F_{sem}(K \cup \{\phi_1, ..., \phi_n\}, \preceq_n) \end{aligned}$$

The second view is not only simpler but more intuitive. More importantly, a revision in this framework is simply an addition. We will adopt the second view in describing our revision semantics in the rest of this paper.

We now reformulate the AGM postulates to suit our new framework. We will denote the revision of K by μ under the F_{sem} semantics as $F_{sem}(K\hat{+}\mu)$, i.e., $F_{sem}(K\hat{+}\mu) = F_{sem}(K \cup \{\mu\})$.

Definition 5.2 Let K be a knowledge set, μ and ν be consistent sentences, and $K\hat{+}\mu$ represent K revised by μ. Then

(R1) $K\hat{+}\mu$ is a knowledge set;

(R2) $F_{sem}(K\hat{+}\mu) \models \mu$;

(R3) $F_{sem}(K) \wedge \mu \models F_{sem}(K\hat{+}\mu)$;

(R4) $F_{sem}(K\hat{+}\mu) \models F_{sem}(K) \wedge \mu$ if $K \wedge \mu$ is consistent;

(R5) $F_{sem}(K\hat{+}\mu)$ is consistent;

(R6) $F_{sem}(K\hat{+}\mu) \equiv F_{sem}(K\hat{+}\nu)$ if $\mu \equiv \nu$;

(R7) $F_{sem}(K\hat{+}\mu) \wedge \nu \models F_{sem}(K\hat{+}(\mu \wedge \nu))$;

(R8) $\quad F_{sem}(K\hat{+}(\mu \wedge \nu)) \equiv F_{sem}((K\hat{+}\mu)\hat{+}\nu)$ if $\mu \wedge \nu$ is consistent. $\Box$

The modification to the AGM postulates is minimum, possibly except R8, in that the postulates are revised only to suit our new framework, and the underlying meanings are not affected. This can be seen from the fact that the revised postulates are exactly the same as the original AGM postulates if both K and $F_{sem}(K)$ are defined as the same deductively closed set of sentences.

In the ideal situations, independence of revision orders is required. That is, K revised with μ first and then ν should be the same as K revised with ν first and then μ, i.e.,

$$(F\hat{+}\mu)\hat{+}\nu \equiv (K\hat{+}\nu)\hat{+}\mu,$$

as long as μ and ν can peacefully live together. The eighth AGM postulate expresses a weaker desire for such independence. Despite the fact that it is weaker than desired, P8 is the main obstacle for many revision semantics to satisfy the AGM postulates [7, 14]. On the other hand, the revised, i.e. R8, implies that

$$F_{sem}((K\hat{+}\mu)\hat{+}\nu) \equiv F_{sem}((K\hat{+}\nu)\hat{+}\mu)$$

that is, a total independence of revision orders. From such a point view, R8 is stronger than P8.

However, there may be different interpretations of the AGM postulates. If P8 is interpreted as

$$(P8') \quad F_{sem}(K\hat{+}(\mu \wedge \nu)) \models F_{sem}(K\hat{+}\mu) \wedge \nu \quad \text{if } F_{sem}(K\hat{+}\mu) \wedge \nu \text{ is consistent,}$$

then R8 is weaker than P8 since $P8'$ implies R8 but not vice versa. We doubt the legitimacy of such an interpretation. Otherwise, it is not difficult to show that no reasonable revision semantics satisfies the postulate, other than either throwing out all old conflicting beliefs or imposing a linear ordering on all beliefs [3, 1].

6 Revision Semantics

In this section, we first define our revision semantics for knowledge sets, and then show that our revision semantics does satisfy the reformulated AGM postulates.

Definition 6.1 Let K be a knowledge set with priority relation $\preceq$, and μ be a new sentence. Then the new knowledge set by revising K with μ is represented by $K\hat{+}\mu = K \cup \{\mu\}$, together with a revised priority relation $\preceq' = \preceq \cup \{\mu \succeq t \mid t \in K\}$. Furthermore, the semantics of K revised by μ is defined as $F_{sem}(K\hat{+}\mu, \preceq')$. $\Box$

The revision of a knowledge set, now, is as simple as an addition, as shown below.

Lemma 6.1 Let K be a knowledge set with priority relation $\preceq$, T_K be the belief theory for K, and μ be a sentence. Then we have

$$F_{sem}(K\hat{+}\mu) \equiv F_{sem}(K \cup \{\mu\}, \preceq) \wedge \mu \equiv MCIR(T_K \wedge \mu; B| \preceq; P). \quad \Box$$

Proof: Let $T_K(P, B)$ and $T_{K_\mu}(P, B) = T_K(P, B) \wedge (\mu \leftarrow B_\mu)$ be two belief theories for K and $K \cup \{\mu\}$ respectively, and $\preceq' = \preceq \cup \{\mu \succeq t | t \in K\}$. Then

$$F_{sem}(K\hat{+}\mu) \equiv MCIR(T_{K_\mu}; B| \preceq'; P) \equiv MCIR(T_K \wedge (\mu \leftarrow B_\mu); B| \preceq'; P).$$

¿From [8], we have

$$\begin{aligned} & CIR(A; P_1 \succ P_2 \succ \ldots \succ P_n; Z) \\ \equiv\ & CIR(A; P_1; P_2, \ldots, P_n, Z) \wedge CIR(A; P_1; P_2 \succ \ldots \succ P_n; Z), \end{aligned}$$

which implies that

$$\begin{aligned} & MCIR(T_K \wedge (\mu \leftarrow B_\mu)); B| \preceq'; P) \\ \equiv\ & MCIR(T_K \wedge (\mu \leftarrow B_\mu); \{B_\mu\}; P \cup (B - \{B_\mu\}) \\ & \wedge MCIR(T_K \wedge (\mu \leftarrow B_\mu); B| \preceq; P) \\ \equiv\ & T_k \wedge \mu \wedge B_\mu \wedge MCIR(T_K \cup \{\mu \leftarrow B_\mu\}; B| \preceq; P) \\ \equiv\ & MCIR(T_K \wedge \mu; B| \preceq; P) \wedge B_\mu \\ \equiv\ & F_{sem}(K \cup \{\mu\}, \preceq) \wedge B_\mu. \end{aligned}$$

Therefore, we have $F_{sem}(K\hat{+}\mu) \equiv F_{sem}(K \cup \{\mu\}, \preceq) \wedge \mu$

(Note that we assume B_μ is also contained in B, though it does not appear in $T_K \wedge \mu$.) $\Box$

Note that, in Definition 6.1, the new priority relation $\preceq'$ means the new sentence μ has the highest priority amongst all sentences in K in the revised knowledge system. This treatment is not necessary at all, and the status of a new sentence can be determined on per applications. We adopt such an approach for the sake of easy comparison. The AGM postulates require that newly added sentences have higher priority to be retained in the revised knowledge system.

The following theorem shows that our revision semantics satisfies the revised AGM postulates.

Theorem 6.1 The revision semantics defined in Definition 5.2 satisfies the revised AGM postulates.

First, we show the following utility lemma.

Lemma 6.2 Assume $A_1(P,Z) \wedge A_2(P,Z)$ is consistent, where A_1 and A_2 are first order theories and P and Z are disjoint sets of predicates in both A_1 and A_2. Then

$$\begin{array}{rl} & MCIR(A_1(P,Z); P| \preceq; Z) \wedge A_2(P,Z) \\ \models & MCIR(A_1(P,Z) \wedge A_2(P,Z); P| \preceq; Z) \end{array}$$

Proof: Assume $MCIR(A_1; P| \preceq; Z) \wedge A_2$ is consistent, otherwise it is trivial. Let $\mathbf{m}$ be a model of $MCIR(A_1; P; Z) \wedge A_2$. Then $\mathbf{m}$ is a model of $A_1 \wedge A_2$, and a *(P, Z)-maximal* model of A_1, i.e., $\mathbf{m}$ is not (P, Z)-smaller than any model of A_1. Furthermore, for any model $\mathbf{n}$ of $A_1 \wedge A_2$, since $\mathbf{n}$ is also a model of A_1, $\mathbf{m}$ is not (P, Z)-smaller than $\mathbf{n}$, i.e., $\mathbf{m}$ is also a (P, Z)-maximal model of $A_1 \wedge A_2$. Therefore, we have

$$MCIR(A_1; P; Z) \wedge A_2 \models MCIR(A_1 \wedge A_2; P; Z)$$

which in turn shows this utility lemma. $\Box$

Now we show the theorem.

Proof of Theorem 6.1 :

Notations: T_K denotes the belief theory of K. T_{K_μ} denotes the belief theory of $K \hat{+} \mu$, i.e., $T_{K_\mu} = T_K \cup \{\mu \leftarrow B_\mu\}$.

B_K denotes the set of belief predicates in T_K, i.e., $B_K = \{B_\mu \mid (\mu \leftarrow B_\mu) \in T_K\}$.

(R1) Trivial.

(R2) It follows from Lemma 6.1.

(R3) Assume $F_{sem}(K) \wedge \mu$ is consistent; otherwise it is trivial. Then, by Lemmas 6.1 and 6.2,

$$F_{sem}(K) \wedge \mu = MCIR(T_K; B| \preceq; P) \wedge \mu \models MCIR(T_K) \wedge \mu; B| \preceq; P).$$

By Lemma 6.1, $MCIR(T_K \wedge \mu; B| \preceq; P)$ is equivalent to $F_{sem}(K \hat{+} \mu)$.

(R4) If $K \wedge \mu$ is consistent, then $F_{sem}(K) = T_K \wedge B_K$ and $F_{sem}(K \hat{+} \mu) = T_K \wedge B_K \wedge \mu \wedge B_\mu$. The postulate obviously holds.

(R5) It is known that a consistent first order theory with only universally quantified variables has at least one (P,Z)-minimal model. This result can be extended to priority circumscription and maximizing circumscription. Clearly, $T_K \cup \{\mu \leftarrow B_\mu\}$ is always consistent even if $K \cup \{\mu\}$ is not.

(R6) It follows from the fact that $T_{K_\mu} \equiv T_{K_\nu}$ if $\mu \equiv \nu$.

(R7) By Lemma 6.1, we have

$$F_{sem}(K\hat{+}\mu) \equiv MCIR(T_K \wedge \mu; B| \preceq; P), \text{and}$$

$$F_{sem}(K\hat{+}(\mu \wedge \nu)) \equiv MCIR(T_K \wedge (\mu \wedge \nu); B| \preceq; P).$$

Assume $MCIR(T_K \wedge \mu; B| \preceq; P) \wedge \nu$ is consistent; otherwise it is trivial. Then, by Lemma 6.2,

$$MCIR(T_K \wedge \mu; B| \preceq; P) \wedge \nu \models MCIR(T_K \wedge (\mu \wedge \nu); B| \preceq; P).$$

(R8) Let $\preceq_\mu = \preceq \cup \{B_\mu \succeq B_t | t \in K\}$. By Lemma 6.1,

$$F_{sem}(K\hat{+}(\mu \wedge \nu)) \equiv MCIR(T_K \wedge \mu \wedge \nu; B| \preceq; P).$$

Since $K\hat{+}\mu$ is $K \cup \{\mu\}$, together with $\preceq_\mu$, we have, by Lemma 6.1,

$$F_{sem}((K\hat{+}\mu)\hat{+}\nu) \equiv MCIR(T_K \wedge (\mu \leftarrow B_\mu) \wedge \nu); B| \preceq_\mu; P).$$

Since B_μ has the highest priority amongst all belief predicates in B with respect to $\preceq_\mu$, and B_μ is consistent with $T_K \wedge (\mu \leftarrow B_\mu) \wedge \nu$, due to the fact that $\mu \wedge \nu$ is consistent, we have

$$MCIR(T_K \wedge (\mu \leftarrow B_\mu) \wedge \nu); B| \preceq_\mu; P) \models \mu \wedge \nu \wedge B_\mu.$$

Therefore,

$$\begin{aligned} & MCIR(T_K \wedge (\mu \leftarrow B_\mu) \wedge \nu); B| \preceq_\mu; P) \\ & \equiv MCIR(T_K \wedge \mu \wedge \nu \wedge B_\mu); B| \preceq_\mu; P) \end{aligned}$$

¿From [8], we have

$$\begin{array}{rl} & MCIR(T_K \wedge \mu \wedge \nu \wedge B_\mu); B| \preceq_\mu; P) \\ \equiv & MCIR(T_K \wedge \mu \wedge \nu \wedge B_\mu; B_\mu; P \cup (B - B_\mu)) \\ & \wedge MCIR(T_K \wedge \mu \wedge \nu \wedge B_\mu); B| \preceq; P). \end{array}$$

However,

$$MCIR(T_K \wedge \mu \wedge \nu \wedge B_\mu; B_\mu; P \cup (B - B_\mu)) \equiv T_K \wedge \mu \wedge \nu \wedge B_\mu.$$

It follows that

$$\begin{array}{rl} & F_{sem}((K\hat{+}\mu)\hat{+}\nu) \\ \equiv & MCIR(T_K \wedge \mu \wedge \nu \wedge B_\mu); B| \preceq; P) \\ \equiv & F_{sem}(K\hat{+}(\mu \wedge \nu)). \quad \Box \end{array}$$

Example 6.1 The following example, which we have used in Example 4.1 to show that the FUV method fails to satisfy the AGM postulates. We now use this same example to show how our approach satisfies the postulates.

$$K = \{a,\ a \leftarrow b \wedge c,\ b,\ c,\ c \leftarrow d,\ d\},$$

with the priority relation $\preceq$:

$$\begin{array}{l} (b) \preceq (a \leftarrow b \wedge c) \\ (c) \preceq (a \leftarrow b \wedge c) \\ (d) \preceq (c \leftarrow d) \preceq (a \leftarrow b \wedge c) \end{array}$$

Let $\mu = \{\neg a\}$ and $\nu = \{\neg d\}$ be the new sentences to be added.

Then K is represented by its belief theory

$$T_K(P, B) = \{a \leftarrow \alpha_1,\ a \leftarrow b \wedge c \wedge \alpha_2,\ b \leftarrow \alpha_3,\ c \leftarrow \alpha_4,\ c \leftarrow d \wedge \alpha_5,\ d \leftarrow \alpha_6\}$$

where

$$\begin{array}{l} P = \{a, b, c, d\} \\ B = \{\alpha_1, \alpha_2, \alpha_3, \alpha_4, \alpha_5, \alpha_6, \alpha_7, \alpha_8, \alpha_9\} \\ \preceq = \{\alpha_2 \preceq \alpha_3,\ \alpha_2 \preceq \alpha_4,\ \alpha_2 \preceq \alpha_5,\ \alpha_5 \preceq \alpha_6\}. \end{array}$$

Now let

$$\begin{array}{l} T_{K_\mu} = T_K \cup \{\mu \leftarrow \alpha_7\} \\ T_{K_{\mu+\nu}} = T_K \cup \{\mu \leftarrow \alpha_7,\ \nu \leftarrow \alpha_8\} \\ T_{K_{(\mu\wedge\nu)}} = T_K \cup \{(\mu \wedge \nu) \leftarrow \alpha_9\} \\ \preceq' = \preceq \cup \{\alpha_i \preceq \alpha_j \mid (1 \leq i \leq 6) \wedge (j = 7, 8, 9)\} \cup \{\alpha_7 \preceq \alpha_8\}. \end{array}$$

Then we have

$$\begin{array}{l} F_{sem}(K \hat{+} \mu) = MCIR(T_{K_\mu}(P, B); B|\preceq'; P) \\ F_{sem}((K \hat{+} \mu) \hat{+} \nu) = MCIR(T_{K_{\mu+\nu}}(P, B); B|\preceq'; P) \\ F_{sem}(K \hat{+} (\mu \wedge \nu)) = MCIR(T_{K_{(\mu\wedge\nu)}}(P, B); B|\preceq'; P) \end{array}$$

It is easy to show that

$$\begin{array}{l} F_{sem}((K \hat{+} \mu) \hat{+} \nu) \equiv F_{sem}(K \hat{+} (\mu \wedge \nu)) \\ F_{sem}(K \hat{+} \mu) \models \{\neg a,\ a \leftarrow b \wedge c,\ c \leftarrow d,\ b \vee c,\ b \vee d\}, \text{ and} \\ F_{sem}((K \hat{+} \mu) \hat{+} \nu) \models \{\neg a,\ \neg d,\ a \leftarrow b \wedge c,\ c \leftarrow d,\ b \vee c\} \end{array}$$

□

7 Final Remarks

The main goal that this paper has achieved is to define a semantics of first order belief revision systems that satisfies the (revised) AGM postulates. The semantics we propose in this paper is based on a form of circumscription, which we have called *maximizing circumscription*. The most important features of this semantics we believe are that (1) a revision operation under this semantics is a simple addition of the new knowledge to the underlying knowledge system, and (2) it deals with belief revision of first order knowledge systems.

The first feature provides a direct implementation strategy for the semantics; its realization is essentially to utililze a circumscription algorithm. The second feature is unique in the literature of belief revision, since the problem of belief revision has so far been dealt with only in the context of propositional language, while in reality a knowledge base is likely to be first order rather than propositional.

References

[1] C. Alchourron, A. Gardenfors, and D. Makinson. On the logic of theory change: Partial meet contradiction and revision functions. *The Journal of Symbolic Logic*, 50:510–531, 1985.

[2] C. Boutilier A logic for revision and subjunctive queries. In *Proc. AAAI-92*, 609-615, 1992.

[3] R. Fagin, J. Ullman, and M. Vardi. On the semantics of updates in databases. In *Proceedings of the 2nd ACM PODS*, pages 352–365, 1983.

[4] G. Grahne, A. Mendelzon, and R. Reiter On the semantics of belief revision systems. *Unpublished manuscript.*

[5] A. Gardenfors and D. Makinson. Revision of knowledge systems using epistemic entrenchment. In *Proceedings of the 2nd Workshop on Theoretical Aspects of Reasoning about Knowledge*. Morgan Kaufmann, 1988.

[6] Hadley, Fagin, and Halpern. On logical omniscience: A critique with an alternative. In *Proceedings of the 6th Canadian Conf. on AI*, pages 49–56, 1986.

[7] H. Katsuno and A.O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52:263 – 294, 1991.

[8] V. Lifschitz. Computing circumscription. In *Proceedings of the 9th Int. Joint Conference on AI*, 1986.

[9] J. McCarthy. Circumscription – a form of non-monotonic reasoning. *AI*, 13:27–39, 1980.

[10] J. McCarthy. Applications of circumscription to formalizing common sense knowledge. *AI*, 28:89–116, 1986.

[11] R.C. Moore. Semantic considerations on non-monotonic logic. *AI*, 25:75–94, 1985.

[12] B. Nebel. Belief revision and default reasoning: Syntax-based approaches. *KR'91*, pages 417–428, 1991.

[13] T. Przymusinski. An Algorithm to Compute Circumscription. In *Artificial Intelligence*, vol 38, pages 47-73, 1989.

[14] L. Willard and L. Yuan. The revised gardenfors postulates and update semantics. In *Proceedings of 1990 International Conference on Database Theory*, pages 409 – 421, 1990.