UNIVERSITY OF ALBERTA

# A New Framework and Update Semantics for Deductive Databases

by



Leigh Willard

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements for the degree of Master of Science.

DEPARTMENT OF COMPUTING SCIENCE

Edmonton, Alberta
Fall 1992

ISBN 0-315-77294-8

UNIVERSITY OF ALBERTA

RELEASE FORM

NAME OF AUTHOR: Leigh Willard
TITLE OF THESIS: A New Framework and Update Semantics for Deductive Databases
DEGREE: Master of Science
YEAR THIS DEGREE GRANTED: 1992



Leigh Willard
#103 10818-81 Avenue
Edmonton, Alberta
T6E 1Y4

Date: 9 October 1992

UNIVERSITY OF ALBERTA

FACULTY OF GRADUATE STUDIES AND RESEARCH

The undersigned certify that they have read, and recomended to the Faculty of Graduate Studies and Research for acceptance a thesis entitled A New Framework and Update Semantics for Deductive Databases submitted by Leigh Willard in partial fulfillment of the requirements for the degree of Master of Science.

Li Yan Yuan

Randy Goebel

Bernard Linsky

Jia You

Date: 2 October 1992

*This thesis is dedicated to my parents,*
*Roy and Bernice Willard,*
*because I couldn't have made it without them.*

## Abstract

This thesis discusses work done on the foundations of deductive database updates, then presents a new approach to revision. Up until this point in time, many update semantics have been proposed, all varying on their definition of minimal change. A set of postulates was given by Alchourron, Gärdenfors, and Makinson to lay the foundation for updates to knowledge bases. Based on that, this work presents a new approach to database updates, partially closed theories. A modified set of postulates is presented, followed by a discussion of the epistemic importance of facts during an update. A binary relation over sentences in a theory is given, and is used to construct an update semantics which satisfies the revised postulates. Finally, some rules are presented for assigning epistemic importance to sentences.

# Acknowledgements

I would like to thank my supervisor Li Yan Yuan for all of his hard work and guidance and for answering all of my questions.

Lots of thanks to Tony Marsland for all of his support and encouragement.

I am grateful to Bernard Linsky, Randy Goebel, and Jia You for being on my committee and for their very useful feedback. As well big thanks to Bill Armstrong for chairing my committee, but especially for giving me encouragement in my first undergraduate year.

Also I cannot say thank you enough to my family for helping me make it this far, and to Tim Breitkreutz. Special thanks goes to Audra Witiuk for her much needed emotional support.

# Contents

# List of Tables

# Chapter 1

# Introduction

A knowledge base needs to be updated as our perception of the world changes. Several kinds of updates may occur. *Revision* is a common type of update; it adds newly acquired knowledge to the system. If the new knowledge does not conflict with current beliefs, the revision is simple – add it to the system. However, if the new knowledge is inconsistent with old beliefs, the conflict must be resolved, usually by derogating some old beliefs. The question is whether to throw away all old beliefs, or to choose a subset of the old beliefs as victims. A criteria is needed for making that decision.

Alchourron, Gärdenfors, and Makinson have proposed a set of rationality postulates (henceforth referred to as the AGM postulates), which are based on philosophical grounds and which establish the foundation for knowledge base revisions [10, 2]. Katsuno and Mendelzon have rephrased the AGM postulates in terms of a model-theoretic point of view, and give a characterization of all revision schemes that satisfy the AGM postulates [14].

Many update semantics have been proposed for knowledge bases and deductive databases [1, 2, 3, 7, 21, 9, 22, 26, 24, 25]. However, as noticed by Katsuno and Mendelzon [13], so far most reasonable update semantics do not satisfy the AGM postulates. Dalal uses the number of propositional letters to choose the victims, and although his method does satisfy the postulates it lacks a rational ground [6, 5, 14, 15]. On the other hand, the partial meet revision semantics, proposed by Alchourron et al., satisfy the AGM postulates by discarding too much information. Although their semantics are not realistic, they do serve as a lower bound for information which should be returned by any revision semantics [2]. Recently a new update semantics has been proposed which is based on a logic calculus and which does satisfy the AGM postulates [4].

We believe that the difficulty in finding update semantics which satisfy the AGM postulates is partly because of the unsuitable framework of the postulates and the ignorance of epistemic importance [7]. Therefore, we propose a new revision semantics for deductive databases by resolving these two problems.

In this dissertation, first the AGM postulates are analyzed in the context of deductive databases. The AGM postulates, though sound on philosophic grounds, are shown to be unrealistic when applied to deductive databases. We argue that modification is necessary, and based on this, *partially closed theories* are proposed as the framework for our deductive database update semantics. Partially closed theories consider the set of minimally derivable sentences. However, unlike some other frameworks, the clauses explicitly expressed in the database shall also be respected. The AGM postulates are modified in the context of partially closed theories. The modification is limited only to the context, the underlying meaning is not changed.

It is generally agreed that epistemic importance of beliefs plays an important role during the update process. That is, those beliefs with higher priority should not be derogated unless it is absolutely necessary [10]. Given a priority relation over the set of beliefs in a knowledge base, a new update semantics is given, which is shown to satisfy the revised AGM postulates.

The question of how to assign epistemic ordering to a set of beliefs is also addressed, with a method for an ordering given. When applied to the update operator, this method yields an intuitive and meaningful result. This update semantics satisfies the AGM postulates in the concerned context.

This dissertation is organized as follows. Chapter 2 recalls some fundamental concepts and states the AGM postulates. It also reviews previous work done on update semantics. The necessity of revising the postulates, partially closed theories, and the revised AGM postulates are presented in Chapter 3. In Chapter 4, we propose a new update semantics and show that the semantics satisfy the revised AGM postulates. Then in Chapter 5 an algorithm for assigning epistemic importance to facts is given, which, when incorporated into the new update semantics satisfies the revised AGM postulates. Chapter 6 provides some examples and comparisons of different update approaches. Chapter 7 gives the conclusion.

# Chapter 2

# Preliminaries

This chapter recalls some preliminary results and concepts of knowledge base updates. The AGM Postulates are given, followed by a survey of work done on the semantics of deductive database updates.

## 2.1 Concepts

The work in this thesis assumes some background in logic. For further information on the theory of logic the reader is directed to consult [17, 18, 8].

Throughout most of this paper, a *Deductive Database* or *Knowledge Base* (also called a logical database) will be taken to mean a set of clauses of the form:

$$a(t) \vee b(t) \vee c(t) \vee \ldots \vee z(t)$$

where $t$ is a set of terms and $a, b, c, \ldots, z$ are atoms or the negation of atoms. For shorthand, the notation $a \vee b \vee \ldots \vee z$ is often used to represent a clause. Any universally quantified formula may be defined in terms of an equivalent set of clauses.

As well, some literature views a database as a closed theory. A closed theory contains all of the logical consequences of a given set of clauses. The operator $Cn$ takes sets of clauses and returns a closed theory. The AGM postulates are proposed in the context of knowledge sets, which are closed theories. The notation $Cn(T)$ or $T^*$ will be used to denote a closed theory.

It will be clear from context which type of database is being referred to.

Deductive databases have many advantages. They have strong expressive power since a single database statement may replace many explicit facts. The database language may be used to represent databases, queries, integrity constraints, views, and programs. Since the theory of logic is well understood, it provides a theoretical background for deductive databases. Also, because of the language, a user is able to expresses databases, queries, integrity constraints and correct answers declaratively, not procedurally. Perhaps one of the most important reasons for accepting deductive databases is because the declarative concept of a correct answer is clear and can be separated from the procedural semantics [17].

Although the meaning of a deductive database is understood, the update semantics are not widely agreed upon.

A database may be either viewed as the collection of clauses, or as one of the interpretations of the database. When viewed as a set of clauses (theory based) an update is applied to the theory. A model based semantics applies an update to the underlying models of the theory.

There are different types of updates to deductive databases, and they may be classified as *revision, contraction, elimination*, and *retraction*. Revision refers to the addition of new facts to a theory, where the new facts may conflict with existing information. The previously held theory must be modified so that the new theory (which includes the new information) is consistent. Contraction is the rejection of a proposition which was previously held to be true. Elimination is the deletion of all knowledge about a specific proposition from the theory. If a fact is to be eliminated, it may be necessary to eliminate other facts from the theory that depend on the fact as well. Retraction refers to the undoing of a previous operation.

It has been argued [10] that the only two meaningful types of updates are contraction and revision. The contents of this dissertation will focus on database revisions. All results may be generalized to deal with contraction as well. There is a direct correspondence between contraction and revision.

Let $\hat{+}$ denote the operation of revision, such that if $T$ is a theory and $\mu$ is a sentence, then $T\hat{+}\mu$ is the revision of $T$ by $\mu$ Similarly, let $-$ denote the operation of contraction. The notation $T \cup \mu$ is the smallest deductively closed set containing both $T$ and $\mu$.

Then revision may be defined in terms of contraction as $T\hat{+}\mu = Cn((T - \neg\mu) \cup \mu)$ (called the Levi identity [10]), and contraction may be defined in terms of revision as $T - \mu = Cn(T \cap (T\hat{+}\neg\mu))$ (called the Gärdenfors identity

[19]).

As there are many different semantics for updating deductive databases, a foundation for updates needs to be established. The following section describes work done to set a standard for knowledge base updates.

## 2.2 AGM Postulates

Let $K$ be a knowledge set, and $\mu$ and $\alpha$ be sentences. The revision of $K$ by $\mu$, denoted $K\hat{+}\mu$, represents a new knowledge set obtained from $K$ by adding new knowledge represented in $\mu$. $K \sim \mu$ is the is the smallest deductively closed set containing both $K$ and $\mu$. The AGM postulates for revision are as follows [10].

**(G1)** $K\hat{+}\mu$ is a knowledge set.

**(G2)** $\mu \in K\hat{+}\mu$.

**(G3)** $K\hat{+}\mu \subseteq K \sim \mu$.

**(G4)** If $\neg\mu \notin K$ then $K \sim \mu \subseteq K\hat{+}\mu$.

**(G5)** $K\hat{+}\mu$ consists of all formulas only if $\mu$ is inconsistent.

**(G6)** If $\mu \equiv \alpha$ then $K\hat{+}\mu = K\hat{+}\alpha$.

**(G7)** $K\hat{+}(\mu \wedge \alpha) \subseteq (K\hat{+}\mu) \sim \alpha$.

**(G8)** If $\neg\alpha \notin (K\hat{+}\mu)$ then $(K\hat{+}\mu) \sim \alpha \subseteq K\hat{+}(\mu \wedge \alpha)$

The first postulate states that the revision of a knowledge set must result in a new knowledge set. The second postulate states that the new knowledge must be retained in the update. Postulate G3 stipulates that the revision of $K$ by $\mu$ is a subset of the knowledge set $K$ appended by $\mu$. In other words, the only information that can be added to the knowledge set is that derived from $K$ and $\mu$. The fourth postulate represents the idea that if K is not inconsistent with $\mu$, the revision is done by simply adding $\mu$ to $K$. If $\mu$ is unsatisfiable, then anything may be derived from $K\hat{+}\mu$, and therefore, should be included in the revision, which is G5. The sixth postulate specifies the principle of irrelevance of syntax, that is, the update of a knowledge set with logically equivalent sentences should give the same result. The seventh states that if the knowledge set $K$ is revised with $\mu \wedge \alpha$, then the revision will be a subset of the knowledge $K \hat{+} \mu$ augmented by $\alpha$. Notice that this postulate means that if $K'$ is the knowledge derogated from K when doing the revision with $\mu \wedge \alpha$, and $K''$ is the knowledge derogated from $K$ when doing the revision with $\mu$, then $K' \supseteq K''$. The last postulate states that if $\alpha$ is consistent with $K \hat{+} \mu$ then $(K \hat{+} \mu) \cup \alpha = K \hat{+} (\mu \wedge \alpha)$ [14]. That is, for the derogated knowledge, $K' = K''$.

The postulates G7 and G8 together represent the idea that the change of $K$ due to the inconsistency of $K$ and $\mu$ must be minimal in the process of revision; no unnecessary derogations are permitted. Katsuno and Mendelzon have defined the minimal change in terms of models of knowledge sets [14].

The postulates stipulate basic rules that are claimed to be necessary for any update semantics. Then given a reasonable update semantics, the postulates should be satisfied. The proof of the postulates will be if this happens.

## 2.3 Update Semantics

This section presents some update semantics to deductive databases. Although a number of different approaches will be outlined, the semantics presented here do not represent all of the existing update semantics.

For clarity, the update semantics are broken down into those which operate on models of a theory, and those which are applied to the theory itself.

Chapter 6 will do a comparison of all the update semantics given against the proposed update semantics.

### 2.3.1 Model Based

Model based semantics are those which consider the effects of an update to the underlying models of a theory. The syntax of a theory is not important, the underlying models are the units of change. If two theories are syntactically different, but have the same models, the result of an update should be the same for both of them.

Since these semantics consider only the models, when an update is done the result should be another model which is "closest" to the original model, and in which the new information is true. The definition of closeness, however, is not generally agreed upon, as illustrated by the different semantics.

### Borgida Semantics

To measure closeness, Borgida considers the number of atoms which differ between two interpretations [3]. Define diff($I$, $\mu$) as a set of all sets of propositional letters on which interpretation $I$ and some model of $\mu$ differ. A set in diff($I$, $\mu$) is minimal if there is no subset of it in diff($I$, $\mu$). Then diff may be used as a measure of closeness between an interpretation and a formula.

When $\mu$ is consistent with the theory $T$, then $T \wedge \mu$ is the result of the update. Otherwise, the models of the resulting theory are chosen from the models of $\mu$. If the model $m$ is chosen to be a model of the updated theory, then there is some model $m'$ of $T$ such that the set of propositional letters on which $m$ and $m'$ differ is minimal in diff($m'$, $\mu$).

The Borgida semantics have been shown [14] to violate the AGM postulate number eight.

**Example 2.1** *Given the database* $T = \{(a \wedge b \wedge c) \vee (\neg a \wedge \neg b \wedge \neg c)\}$*, revise it with the new knowledge* $\mu = \{(a \wedge b \wedge c) \vee (a \wedge \neg b \wedge \neg c)\}$*. There are two models of* $T$*,* $T_1 = \{a, b, c\}$*, and* $T_2 = \emptyset$*. There are two models of* $\mu$*,* $\mu_1 = \{a, b, c\}$ *and* $\mu_2 = \{a\}$*.* $Diff(T_1, \mu) = \{\emptyset, \{b, c\}\}$*, and* $Diff(T_2, \mu) = \{\{a, b, c\}, \{a\}\}$*. Since the difference between* $\mu_1$ *and* $T_1$ *is minimal in* $Diff(T_1, \mu)$*, then* $\mu_1$ *is included in the models of* $T + \mu$*. Since the difference between* $\mu_2$ *and* $T_2$ *is minimal in* $Diff(T_2, \mu)$*, then* $\mu_2$ *is also included in the models of* $T + \mu$*. Then* $T + \mu$ *is a theory with models* $\{\mu_1\}$ *and* $\{\mu_2\}$*, so* $T + \mu = \mu$*.* □

### Dalal Semantics

Dalal defines minimal change by the number of propositional letters on which two interpretations differ. Some notation is needed to formally represent the Dalal revision [6].

Let $w$ be an interpretation over a set of atoms $\Lambda$, then let g($w$) be the set of all interpretations that differ from $w$ in at most one atom in $\Lambda$ (notice that $w$ is in this set). When $A$ is a set of interpretations, let $g(A) = \cup_{w \in A} g(w)$. If $\Psi$ is a formula then define $G(\Psi)$ by its set of models as $mod(G(\Psi)) = g(mod(\Psi))$. Let $g_i(A)$ be defined in the usual way as $g_0(A) = A$, and $g_i(A) = g_{i-1}(g(A)))$ otherwise.

Dalal's revision operator is defined as $\Psi + \mu = G_k(\Psi) \cup \mu$, where $k$ is the least value of $i$ for which $G_i(\Psi) \cup \mu$ is consistent.

The revision may be restated so that the result of an update of $\mu$ to $T$ is any formula which has $M'$ as a set of models where

$\forall m \in M'$

- $m \models \mu$
- no other model satisfies the first point and differs from any model of $T$ by fewer atom than $m$ differs from the model of $T$.

The Dalal revision has been shown to satisfy the AGM Postulates [5].

**Example 2.2** *Consider again example 2.1. Since the $diff$ operator of Borgida may also be used to compute the Dalal semantics, it will be used here, letting the result of a $diff$ operation be the number of propositional letters which differ, rather than the set of propositional letters which differ. Then $diff(\mu_1, T_1) = 0$, $diff(\mu_1, T_2) = 3$, $diff(\mu_2, T_1) = 2$, and $diff(\mu_2, T_2) = 1$. Since the distance between $Mod(T)$ and $\mu_1$ is 0, and the distance between $Mod(T)$ and $\mu_2$ is 1, then $\mu_1$ is the result of the update.* □

### Jackson and Pais

The update philosophy here is that given a view of the world, when presented with a new piece of information we should try to find the closest world which is consistent with the new information [21].

The semantic revision of the theory $T$ by $\mu$ (denoted $\mu(Mod(T))$ is given by:

- $\emptyset$ if $Mod(T) = \emptyset$
- If $Mod(T) \cap Mod(\mu) \neq \emptyset$ then $\mu(Mod(T)) = Mod(T) \cap Mod(\mu)$ else $\mu(Mod(T)) = closest(Mod(\mu), Mod(T))$

Where $closest(A, B)$ is an operator which chooses the closest worlds in $A$ to $B$ as follows: $closest(A, B) = \{a \in A \mid (\exists b \in B) N(a, b)\}$, and $N(a, b)$ is a relation representing neighbor where $a$ is a neighbor of $b$ if and only if $a$ is a greatest lower bound (glb) or least upper bound (lub) of $b$, or there is no other world in $A$ which is a glb or lub of $b$, and there is no other world in $A$ that differs from $b$ on fewer atoms than $a$ under set inclusion.

These semantics violate postulates K5 and K8.

**Example 2.3** *Use again example 2.1. Since* $Mod(T) \cap Mod(\mu) \neq \emptyset$*, then the result of the update is the intersection of the models, which is* $\{a \wedge b \wedge c\}$*.* □

### Satoh Semantics

Satoh has an approach to revision which divides information into two categories: knowledge and beliefs [22]. Knowledge are facts and cannot change once they are known (hence they are monotonic). Conversely, beliefs are guessed to be true but cannot be taken as a fact, and may change. Belief is nonmonotonic. Satoh's update operator is based on minimal belief revision.

To be able to compare Satoh's update operator with other work, the set of knowledge will be taken to be empty, and all that is known about the world will be considered a belief. The examples will be limited to propositional logic although the method is applicable to first order logic.

The difference is computed between all models of the knowledge set and all models of the belief set, and pairs of knowledge/belief sets with minimal differences are chosen.

Minimal belief revision is then described as: $Mod(T + \mu) = \{\alpha \mid \alpha \in Mod(\mu) \wedge \exists m \in Mod(T)(diff(m, \mu)$ is minimal (set inclusion) in $Diff(T, \mu))\}$

Where $diff(m_1, m_2)$ is the set of propositions on which models $m_1$ and $m_2$ differ and $Diff(T, \mu) = \cup_{I \in Mod(T)} Diff(I, \mu)$, and $Diff(I, \mu)$ is the collection of all sets of propositional letters on which $I$ and some model of $\mu$ differ.

Satoh's semantics do not satisfy the AGM postulates [5].

**Example 2.4** *Consider again example 2.1. Then* $Diff(T, \mu) = \{\{\}, \{b, c\}, \{a, b, c\}, \{a\}\}$*. The difference between* $\mu_1$ *and* $T_1$ *is minimal in* $Diff(T, \mu)$*, so it is one of the models of* $T + \mu$*. However, the difference between* $\mu_2$ *and* $T_1$ *is not minimal in* $Diff(T, \mu)$*, and neither is the difference between* $\mu_2$ *and* $T_2$*, so* $\mu_2$ *is not one of the models of* $T + \mu$*. Then* $T + \mu = \{a \wedge b \wedge c\}$*.* □

### Winslett Semantics

In [25] Winslett describes the update semantics of the possible models approach (PMA) of incorporating new information into logical theories. Her work is based on the possible worlds approach given by Ginsberg [12], which is theory based. The possible states of the world are the models of the theory, and the PMA considers the effect of an action on each possible state of the world. The meaning of applying an action on a theory is the union of the update on all models of the theory.

More formally, let $T$ be a theory, and $T'$ those formulas of $T$ which always must hold. Let $\Lambda$ be a set of sentences to update the theory with. Then the result of the update is a set of models $\cup_{m \in Mod(T)} Incorporate(\Lambda, m)$ where $Incorporate(\Lambda, m)$ is the set M of models such that:

- $M \models \Lambda$ and $M \models T'$
- no other model satisfies the first point and differs from $m$ by fewer atoms (using set inclusion) than $M$ differs from $m$.

Winslett's semantics violate the AGM postulates because when a theory is updated with new knowledge that is consistent with the old, the result may not be equivalent to the conjunction of the new and old knowledge [5].

**Example 2.5** *Let the theory* $T = (a \wedge \neg b) \vee (\neg a \wedge b)$, $T' = \emptyset$, *and let the update* $\mu = \{\neg b\}$. *There are two models of* $T$, $m_1 = \{a\}$, *and* $m_2 = \{b\}$. *The two models are changed to make* $\mu$ *true, so the models become* $m_1' = \{a\}$, *and* $m_2' = \emptyset$. *The resulting theory has as models* $m_1'$ *and* $m_2'$, *so the result of the update* $T + \mu = \{\neg b\}$. □

## 2.3.2 Theory Based

Theory based semantics are concerned with treating a database as a collection of sentences. When an update is done, it is done on the theory itself, as opposed to the models of the theory.

The advantage of theory based semantics over models based is that they can preserve rules. Also, because they are dealing with formulas, they can place importance on some formulas over others. Model theoretics has no mechanism for doing so.

## Alchourron, Gärdenfors, and Makinson Semantics (AGM)

This work deals with contracting a proposition from a closed theory, and revising a closed theory with a new proposition.

The contraction operator, denoted as $-$ and called choice contraction or maxichoice contraction, is defined as follows. Let $T \rfloor \mu$ be the set of all maximal subsets of $T$ that fail to imply $\mu$. The operation $T - \mu$ may then be defined as $\cap(T \rfloor \mu)$ when $T \rfloor \mu$ is nonempty, and as $T$ when $T \rfloor \mu$ is empty. This set in general is far too small. When $T$ is a theory with $\mu \in T$ then $T - \mu = T \cap Cn(\neg\mu)$. In other words, the only things left after the contraction of $\mu$ from $T$ are the logical consequences of $\neg\mu$ that were in $T$ to begin with.

Revision, termed maxichoice revision, is defined similar to contraction using the Levi identity as $T + \mu = Cn((T - \neg\mu) \cup \mu)$. If $\mu$ is inconsistent with $T$ then $T + \mu$ reduces to $Cn((T \cap Cn(\mu) \cup \mu) = Cn(\mu)$. In other words, if a new fact to be added to a theory that is inconsistent with it, then the theory is thrown away, and the resulting theory is the closure of the new fact.

Because the choice contraction function yields so many subsets, there is interest in defining operators that yield fewer subsets. *Partial meet contraction functions* were defined so as to meet this goal. Let $F$ be a function that picks out a class of *most important* subsets of $T$ that fail to imply $\mu$, that is, $F(T \rfloor \mu)$ is a nonempty subset of $T \rfloor \mu$ if $T \rfloor \mu$ is nonempty, and is $\mu$ in the limiting case when $T \rfloor \mu$ is empty. Then the contraction operation $-$ can be defined as $T - \mu = \cap F(T \rfloor \mu)$ for all $\mu$, and is called the partial meet contraction over $T$ determined by $F$. The corresponding revision function is defined using the Levi identity as $T + \mu = Cn((T - \neg\mu) \cup \mu)$.

The *partial meet contraction/revision function*, proposed by Alchourron et al. [2] [19], is an update semantics that satisfies the AGM postulates.

**Example 2.6** *Let $T = \{a, a \rightarrow b\}$, and let $\mu = \neg b$. The closed set of $T$, denoted as $T^* = \{a, a \rightarrow b, b, a \wedge b, a \vee b, b \rightarrow a, \ldots\}^*$. There are two maximal theories consistent with $\mu$, which are $\{\neg b, a, a \vee b, b \rightarrow a\}^* = \{\neg b, a\}^*$, and $T_2^* = \{\neg b, a \rightarrow b, b \rightarrow a\}^* = \{\neg b, a \rightarrow b\}^*$. The intersection of these two theories is $\{\neg b\}^*$, which is the result of the update.* □

## Fagin, Ullman, and Vardi Semantics (FUV)

As well as the update semantics which FUV have proposed, they have two important contributions:

1. they highlight the problems related to closed theories
2. they consider the priority relation among all sentences in a theory.

These two points will be discussed further in the next chapters. Here we will give their semantics as defined on an arbitrary theory (without priorities).

The principles which underlie the FUV update semantics are [7]: when a theory is updated, the resulting theory should have minimal change from the original one, and if there is no one theory with minimal change, then the new theory should reflect that there is more than one theory which accomplishes the update minimally.

When there is more than one theory that updates $T$ with minimal change, then the world must be represented by the models of all the theories that are minimally changed from the original. If $T'$ is the new theory that accomplishes the update of $T$, then $\text{Mod}(T') = \cup_{i \geq 1} \text{Mod}(T_i')$, where each $T_i'$ is a theory that accomplishes the update of $T$ with minimal change. It has been established that this class is well defined when there are only a finite number of theories that accomplish the update minimally.

Let $T_1, \ldots, T_n$ be theories, and let $T'$ be the theory $\vee_{i=1 \ldots n} T_i$, and let $T'' = \cup_{i=1 \ldots n} T_i^*$. Then $\text{Mod}(T') = \text{Mod}(T'') = \cup_{i=1 \ldots n} \text{Mod}(T_i)$. Therefore if $T_1, \ldots, T_n$ are theories that accomplish an update minimally, then $T' = T_1 \vee \ldots \vee T_n$ is the result of the update.

In [7] the semantics were also given of deleting $\mu$ from closed theory $T^*$ as $(\neg\mu \vee T)^*$ when $\mu \in T^*$, and when inserting $\mu$ into $T$ as $\{\mu\}^*$ when $\mu \notin T^*$. Since this results in information loss when $\mu \notin T^*$, they chose not to consider closed theories.

As indicated in [14], the FUV semantics may also discard all old knowledge should inconsistency occur when the closed set is considered.

The FUV semantics is syntax dependent and therefore does not satisfy the AGM postulates.

**Example 2.7** *Let $T = \{a, a \rightarrow b\}$ and let the update $\mu = \neg b$. There are two minimal subsets of $T$ which are the results of updating the theory $T$ with $\mu$, $T_1' = \{a\}$, and $T_2' = \{a \rightarrow b\}$. Then $T_1' \vee T_2' = \{\}$, and the result of the update $T + \mu = \mu = \{\neg b\}$.* □

### Kuper, Ullman, and Vardi Semantics (KUV)

One of the principles of updating a database by KUV is that when inserting or deleting a fact, as little information as possible (explicitly stated facts) should be deleted from the original database [9]. A theory that meets this criteria is a candidate for the resulting database. When there is only one theory that meets the requirement, then the update is straightforward. However, there may be more than one theory that has minimal change from the original database. Therefore, [9] define the result of an update as a set of theories, called a *flock*. The database is a model of one of the theories in the flock.

To update a flock, each theory in it must be updated, and all of the new theories become a theory in the resulting flock.

Formally, the update operation will be defined. Let $S = \{S_1, \ldots, S_n\}$ be a flock. A flock $Q_j = \{T_1, \ldots, T_n\}$ accomplishes the update of $S$ minimally if $T_i$ accomplishes the update of $S_i$ minimally, $1 \leq i \leq n$. Since there could be more than one such flock $Q_j$, the result of the update should be a flock $Q$ such that $Mod(Q) = \cup_{1 \leq j \leq n} Mod(Q_j)$.

The result of an update on a flock $S = \{S_1, \ldots, S_n\}$ can then be defined as a flock $S' = \{S_i^k \mid 1 \leq i \leq n, 1 \leq k \leq m\}$, where $S_i^1, \ldots, S_i^m$ are theories that accomplish the update of $S_i$ minimally.

**Example 2.8** *Let* $T = \{\{a, a \rightarrow b, a \rightarrow c\}\}$ *be a flock, and update it first with* $\mu_1 = \neg b$. *Then the resulting flock is* $T' = \{\{a, \neg b, a \rightarrow c\}, \{a \rightarrow b, \neg b, a \rightarrow c\}\}$. *Update this flock with* $\mu_2 = \neg c$. *The result of the second update is a flock* $T'' = \{\{a, \neg b, \neg c\}, \{a \rightarrow c, \neg b, \neg c\}, \{a \rightarrow b, \neg b, a \rightarrow c\}\}$. □

# Chapter 3

# Revised AGM Postulates for Deductive Database Updates

The necessity to revise the AGM postulates arises from their applicability to deductive databases. It will be shown that the AGM postulates may not be realistic in the deductive database setting. Following is an example which illustrates a revision that satisfies the AGM Postulates.

**Example 3.1** *Given the database* $\{a, b\}$*, revise it with the fact* $\neg a$*. To compute the update, first calculate the closure of the set* $\{a, b\}^*$*. The result is a set* $\{a, b, a \vee b, a \wedge b, a \rightarrow b, b \rightarrow a, \ldots\}$*. Revise this with* $\neg a$ *and there are two minimal results:* $\{\neg a, b\}^*$ *and* $\{\neg a, \neg b\}^*$*. Note that* $\{\neg a, \neg b\}^*$ *is obtained because of the fact that* $\{a \rightarrow b, b \rightarrow a, \ldots\}$ *is a maximal subset of* $\{a, b\}^*$ *that is consistent with* $\neg a$*.* □

Intuitively, if we have the facts "Lemons are yellow" and "Oranges are blue" and then revise the database with the fact "Oranges are not blue", we still expect the fact "Lemons are yellow" to be in the database. If our original database under closure is considered, then two theories result, one in which lemons are yellow and one in which they are not. Since both theories undergo minimal changes, it is not possible to choose between them without further information. Only one resulting theory should be accepted (the one in which "Lemons are yellow" is true), but using the AGM postulates in the context of the closed theory both theories are equally valid.

The reason that undesirable results are produced is that when the closure of a theory is calculated, anything that is a logical consequence of the theory is included in the closure (logical omniscience), and the logical consequences are 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 [16].

The closed theory in essence forces independent pieces of information to "marry". For example, given the facts "Lemons are yellow" and "Oranges are blue", under closure the fact "Lemons are yellow $\wedge$ Oranges are blue" will be added to the theory. Once a set of formulas is closed under deduction, relationships are formed between the various facts, but this is not always desirable.

Additionally, there are many sets of facts that may not be represented by a closed theory. For example, there is no closed theory that contains only the facts "Lemons are yellow" and "Oranges are blue". Clearly this is unacceptable.

The end user of a database cannot be expected to work with knowledge sets. People do not compute relationships between facts when given seemingly independent information. Similarly, databases should not infer relationships between facts.

David Makinson (one of the founders of the AGM postulates) claims that when humans revise their belief set they do it on a finite, or at least recursively enumerable base of a theory, not on a set closed under Cn[19]. This base will in general be either irredundant or reasonably close to irredundant. When the maxichoice operator is applied to a finite base it does not yield an inflated set.

A more serious problem with closed theories is what happens under update. A theorem from FUV shall be restated here [7]. Let $T$ be a closed theory, and let $a$ be a sentence. If $a \notin T$ then the closed result of inserting $a$ into $T$ is $\{a\}^*$. Closed theories suffer from major information loss under update. To see this look at example 3.1. Adding $\neg a$ to the database $\{a, b\}^*$ yields two minimal results, and the conjunction of those results is $\{\neg a\}^*$

The criticisms of closed theories are quite severe. Since the AGM Postulates are written in the context of closed theories, it appears necessary to question the relevance of the AGM postulates.

In a paper by Katsuno & Mendelzon [15] it is argued that there are two types of modifications to a database, and the AGM postulates are only applicable to one type. The first they call *revision* which is used to obtain information about a static world. For example, if a database contains the facts $\{((c \wedge s) \vee (f \wedge a)), \neg s \rightarrow a\}$ and it is revised with $\{\neg s\}$, then $\{f \wedge a\}$ should be deduced.

The second type of modification they termed *update*, which means bringing the knowledge base up to date when the world that it is describing changes. Using the same scenario, $f \wedge a$ could not be deduced. To see why, let $c$ represent the fact that Fido is on the couch and $s$ be the fact that Fido is sleeping. $f$ means that Fido is on the floor, and $a$ means that Fido is eating. Since Fido is awfully quiet, we assume that he is either on the floor eating, or on the couch sleeping. So we yell at him to wake him up ($\neg s$). After this, it is not correct to assume that Fido is on the couch. All that can be determined is that he is not sleeping.

Katsuno and Mendelzon claim that the postulates are not applicable to updates. To see this, consider the work of Dalal, which satisfies all of the postulates. The Dalal operator works by considering all of the models of a theory and the models of the update. The result is the models of the update which are closest to the models of the original theory. The other models of the update are thrown away. This is consistent with a static view of the world where we are narrowing down the unsure information, but is not appropriate in a changing world where all of the previous models should be changed to reflect a change in the world.

The question that the criticisms of the closed theory and AGM Postulates bring up is what do we mean by a set of sentences? Only when we can decide what is meant by a set of sentences can we possibly hope to come up with a reasonable update semantics.

Clearly when humans have a theory such as $\{a, a \rightarrow b\}$, they do not include facts such as $b \rightarrow a$ as part of that theory. Although all logical consequences should not be included in a theory, it will be argued that minimally derived consequences should be.

The model theoretical work on deductive databases applies updates on the models of a theory because they believe that the models are the intended meaning of a theory and that changing the models in face of an update should give the desired results. However, we must recognize that the sentences of a theory must be preserved as well.

Suppose we have as a theory the knowledge that "Socrates is a person" ($p$), and "Socrates is mortal" ($m$), and also "If Socrates is a person, then Socrates is mortal" ($p \rightarrow m$) (example from [12]). Three world descriptions of this could be:

- $\{p, m\}$
- $\{p, p \rightarrow m\}$
- $\{p, p \rightarrow m, m\}$

All worlds are logically equivalent. If it is determined that Socrates is not a person, the first world still would show that he is mortal. The second world would not infer that Socrates is mortal, and the third world would contain two possibilities, encompassing both of the other approaches.

The world containing the minimally derived facts does indeed seem to be the intended meaning because it preserves enough information to derive the other two worlds, but it does not contain all of the information of a closed theory.

The next example also shows that ignoring derived facts may be undesirable.

**Example 3.2** *Consider the statements "A dog with normal legs can run" and "A dog who can run can walk": $\{l \rightarrow r, r \rightarrow w\}$. Add the fact "Fido has normal legs": $\{l\}$. Then modify the database with the fact that "Fido can not run": $\{\neg r\}$. Consider only the given set of clauses. We need to perform $T \dot{+} \mu$, where $T = \{l, l \rightarrow r, r \rightarrow w\}$ and $\mu = \neg r$. There are two maximal subsets of $T$ that are consistent with $\mu$, that is, $\{l, r \rightarrow w\}$ and $\{l \rightarrow r, r \rightarrow w\}$. Naturally, the fact that Fido can not run may imply that he does not have normal legs, but it does not mean he can't walk. However, $w$ can not be inferred from either subset of $T$ above, though $w$ is a logic consequence of $T$. Let $T'$ be $\{l, l \rightarrow r, r \rightarrow w, r, w\}$. That is, $T'$ is obtained from $T$ by adding all minimally derived clauses of $T$. Then $T'$ has two maximal subsets that are consistent with $\mu$, i.e., $\{l, r \rightarrow w, w\}$ and $\{l \rightarrow r, r \rightarrow w, w\}$. We may choose one of these two subsets as the maximal set to carry out the updates, as discussed below. Regardless of the choice, the fact that Fido can walk will be preserved.* □

Obviously, the concerned framework is critical in dealing with update semantics. Some update semantics are applied to deductively closed sets such as the AGM postulates (such as the AGM semantics), some are defined in terms of the theory itself (the FUV semantics), while others may consider the model-theoretical point of view which totally ignores the difference between facts and inference rules (the Winslett semantics). As shown by the above examples and the following examples, those frameworks may not be suitable for deductive database update semantics. For these reasons, partially closed theories are introduced. It will be shown that that within the context of a partially closed theory, a more reasonable update semantics for deductive databases may be found.

**Definition 3.1** *Let $T$ be a theory. The* partially closed theory *of $T$, denoted as $T^+$, is defined as a theory $T^+ = T \cup \{\mu \mid T \models \mu \,\&\, T \not\models \mu' \text{ for any } \mu' \subset \mu\}$.* □

A partially closed theory is essentially a set of minimally derived clauses from the given theory. However, all clauses in the original theory are considered as having higher priority than others, and therefore, are included in the partially closed theory. A theory is said to be a partially closed theory if it is the partially closed theory of itself. A clause $\alpha$ in $T$ is said to be *underlined* if $T \models \alpha'$ for some $\alpha' \subset \alpha$.

Obviously all underlined clauses of $T^+$ are contained in $T$. The difference between the partially closed theory and the theory itself lies in the derived clauses; the difference between the partially closed theory and the set of all minimally derived clauses lies in the underlined clauses; and the difference between the partially closed theory and the knowledge set is that the knowledge set contains all possible underlined clauses.

**Example 3.3** *Let $T$ be $\{a, a \rightarrow b\}$. Then $T^+ = \{a, a \rightarrow b, b\}$ is the partially closed theory of $T$, but $T$ itself is not partially closed ($a \rightarrow b$ is an underlined clause). Furthermore, $T^*$, the knowledge set of $T$, is $\{a, b, a \rightarrow b, b \rightarrow a, a \vee b, a \wedge b, \ldots\}$, and the minimally derived set of $T$ is $\{a, b\}$.* □

Now the AGM postulates will be revised in the context of partially closed theories. Let $T$ be a partially closed theory and let $\Lambda$ and $\Omega$ be sets of clauses. $T \hat{+} \Lambda$ refers to the partially closed theory that is the revision of $T$ by $\Lambda$.

The revised postulates are proposed to overcome the problems of knowledge sets.

**Definition 3.2** *The following are the* Revised AGM Postulates *in the context of partially closed theories* [1].

[1]the postulates have been relaxed from the conditions given in [23], although the meaning is the same.

1. $T \hat{+} \Lambda$ *is a partially closed theory.*

2. $\Lambda \in (T \hat{+} \Lambda)$.

3. $(T \cup \Lambda) \models (T \hat{+} \Lambda)$

4. $(T \hat{+} \Lambda) \models (T \cup \Lambda)$ *if* $T \cup \Lambda$ *is consistent.*

5. $T \hat{+} \Lambda$ *consists of all the propositional formulas only if* $\Lambda$ *is inconsistent.*

6. $(T \hat{+} \Lambda) = (T \hat{+} \Omega)$ *if* $\Lambda^{+} = \Omega^{+}$.

7. $(T \hat{+} \Lambda) \cup \Omega \models T \hat{+} (\Lambda \cup \Omega)$

8. $T \hat{+} (\Lambda \cup \Omega) \models (T \hat{+} \Lambda) \cup \Omega$ *if* $(T \hat{+} \Lambda) \wedge \Omega$ *is consistent.*

The changes to the AGM postulates are minimal in that they have only been changed to operate on partially closed theories. The underlying meaning of the postulates has not changed. The change to postulate six indicates that syntax is no longer absolutely irrelevant for updates.

# Chapter 4

# Meaningful Revision

This chapter presents a semantic characterization of an update, based on the concept of partially closed theories.

A partially closed theory has been shown to mimic the real world by encapsulating the rules and facts that make up the world. Given a desirable framework for knowledge bases, updates are still not straightforward because if a new piece of information conflicts with the knowledge of the world, then the question still arises as to which information should be preserved and which information should be derogated from the set of beliefs.

The following example shows that using the foundation of partially closed theories, an update may still have more than one possible result.

**Example 4.1** *Assume that $\kappa$ is a predicate representing the fact that an apple is mature and disconnected from the tree, and $\epsilon$ is a predicate representing the fact that the apple is on the ground. Then by Newton's law of gravity, we have $\kappa \rightarrow \epsilon$. Let $T$ be a database that represents our belief that the apple is mature and disconnected from the tree, that is, $T = \{\kappa, \kappa \rightarrow \epsilon\}$. However, later we find out that the apple is not on the ground, so $\neg\epsilon$ is observed. To revise the system, we first take the partially closed theory of $T$, $T^+ = \{\kappa, \kappa \rightarrow \epsilon, \epsilon\}$. Then there are two theories that may achieve the update with $\neg\epsilon$, $\{\kappa, \neg\epsilon\}$ and $\{\kappa \rightarrow \epsilon, \neg\epsilon\}$* □

By a human reasoning pattern, we usually do not suspect Newton's law of being in error, but we would suspect that the apple may not be mature and is still on the tree. Therefore, we would place the rule $\kappa \rightarrow \epsilon$ as being more important than than the fact $\kappa$. Thus, the revised theory $\{\kappa \rightarrow \epsilon, \neg\epsilon\}$ should be adopted as the revision of $T^+$ by $\neg\epsilon$.

The example illustrates that some sentences have higher priority to survive after an update, when a conflict occurs. That priority is termed *epistemic importance* [20].

A binary relation (augmented with the obvious transitive rule) may be used to represent or specify epistemic importance amongst all sentences in a theory. That is, given a partially closed theory $T$ and a binary relation $\geq$ on $T$ we say that $s_i \in T$ is epistemically at least as important as $s_j \in T$ if $s_i \geq s_j$.

The binary relation $\geq$ is used to assign importance to sentences in a theory. If sentence $s_i \geq s_j$ then $s_i$ has at least as high priority to survive derogation of a theory as $s_j$ does.

**Example 4.2** *In the preceding example, we would say that* $\kappa \rightarrow \epsilon \geq \kappa$. *This gives* $\kappa \rightarrow \epsilon$ *higher priority than* $\kappa$, *since we feel that Newton's law is more valid than our observations.* □

It is useful to be able to assign an ordering to sentences in a theory such that the elements in the theory form a sequence. The next definition describes such a sequence.

**Definition 4.1** *Let* $T$ *be a partially closed theory and* $\geq$ *a binary relation on* $T$ *representing epistemic importance in* $T$. *Then a sequence* $s_1, \ldots, s_n$ *of all sentences in* $T$ *is said to be an* epistemic sequence *of* $T$ *with respect to* $\geq$ *if when* $s_i \geq s_j$ *and* $s_j \not\geq s_i$ *then* $i < j$. □

Notice that the binary relation is not total. It is not even a partial order, since it is possible that $s_i \geq s_j$ and $s_j \geq s_i$ with $s_i \neq s_j$ (although the priority of $s_i$ equals the priority of $s_j$). We can say nothing from the definition about whether $s_j \geq s_i$ when $j < i$.

Given a partially closed theory there may be more than one epistemic sequence. It is possible to have $s_i \geq s_j$ and also $s_j \geq s_i$. In this case they have similar priority and from the definition of an epistemic sequence either one may be selected ahead of the other.

**Example 4.3** *In the apple/Newton's law example, with the theory* $\{\kappa, \epsilon, \kappa \rightarrow \epsilon\}$, *if we have that* $\kappa \rightarrow \epsilon \geq \kappa$, $\kappa \rightarrow \epsilon \geq \epsilon$, $\epsilon \geq \kappa$, *and* $\kappa \geq \epsilon$ *we can give the* *epistemic sequence as either* $\kappa \rightarrow \epsilon, \kappa, \epsilon$ *or* $\kappa \rightarrow \epsilon, \epsilon, \kappa$. *In either ordering, Newton's law has highest priority.* □

Here it is worthwhile to mention the work done in [7]. Fagin, et. al. considered a logical database to be a theory with a priority attached to each sentence. Since their priorities are given by natural numbers, their priorities form a total ordering (the relation $\geq$ proposed here is not). Revision is accomplished by considering each group of sentences with the same priority. The update is done on each group in turn. They do not discuss how priorities should be assigned to sentences.

In the work reviewed on counterfactual knowledge it was seen that various model based update approaches were proposed, all of them differing on their priority for selecting models of the update. For example, in Winslett's approach, priority is placed on models of the update which are closest to the models of the world, where closeness is measured by set inclusion. Only those models which are closest are considered to be the result of the update. In contrast, here we are placing priority on the sentences of the theory instead of on the models of the theory. Since the foundation for the revised postulates is based on the theory, and not the models of the theory, placing priority on sentences seems to be more realistic.

Given an epistemic sequence, it is simple to incorporate update sentences into it, since the ordering states which sentences have higher priority to survive derogation if a conflict occurs. The following definition gives the straightforward semantics for incorporating new sentences into an epistemic sequence.

**Definition 4.2** *Let* $T$ *be a partially closed theory and* $\geq$ *be a binary relation over the sentences in* $T$. *Then a subset* $T'$ *of* $T$ *is said to be an* epistemically consistent subset *of* $T$ *with sentences* $\mu$ *if there exists an epistemic sequence* $s_1, \ldots, s_n$ *of* $T$ *with respect to* $\geq$ *such that for each* $s_i \in T - T'$, *with* $s_{i_\geq} = \{s_k \mid s_k \in T' \ \& \ k < i\}$ *we have:*

- $s_{i_\geq} \wedge \mu$ *is consistent*
- $s_{i_\geq} \wedge \mu \wedge s_i$ *is inconsistent*

□

The notation $sub(S, \mu)$ will be used to denote the epistemically consistent subset given the epistemic sequence $S$ and update sentences $\mu$.

The definition states that every epistemically consistent subset $T'$ has an associated epistemic sequence which is used to select the elements of $T'$ according to epistemic importance. Every sentence in the sequence is considered in turn (and hence in order). If the candidate sentence from the sequence is consistent with the update sentences, and is consistent with the sentences chosen from the theory so far, it is accepted. Otherwise it is rejected. This strategy proceeds for each sentence in the sequence. The sentences which come after the candidate sentence are not considered when making a decision to reject or accept the sentence.

The next theorem is an important result for epistemically consistent subsets.

**Observation 1** *Given an epistemic sequence of $T$ w.r.t. $\geq$ and update sentences $\mu$, there is only one epistemically consistent subset.* □

*Proof:* Assume the contrary, that there is more than one subset. Let $T_1$ and $T_2$ be two epistemically consistent subsets, and the epistemic sequence of $T$ w.r.t. $\geq$ is $s_1, s_2, \ldots, s_n$.

Then let

$$s_i \in T_1 - T_2 \textit{ such that } \forall k < i((s_k \in T_1, s_k \in T_2) \vee (s_k \notin T_1, s_k \notin T_2))$$

Then $s_i$ is the first element in the sequence which is in $T_1$ but not $T_2$. Since $s_i \in T - T_2$ we have that $s_{i_\geq} \wedge \mu$ is consistent and $s_{i_\geq} \wedge \mu \wedge s_i$ is inconsistent. But $s_{i_\geq} \in T_1$ so $s_{i_\geq} \wedge \mu \wedge s_i$ must be consistent. Contradiction. □

Thus an epistemically consistent subset is based on intuition. Given a partially closed theory, the higher priority sentences are more likely to be in the subset than lower priority sentences. Given an epistemic sequence it is then easy to add update information. The next example shows how to obtain an epistemically consistent subset.

**Example 4.4** *Take the apple/Newton's law example from this chapter. An epistemic sequence may be $\kappa \rightarrow \epsilon, \kappa, \epsilon$. Then we select the maximal subset of $T^+$ consistent with $\neg\epsilon$ and the sequence. Since $\kappa \rightarrow \epsilon$ has no lower epistemic importance than $\kappa$, it is first considered to be in the consistent subset. Since* *it is consistent with* $\neg\epsilon$ *it is added to the result. Then* $\kappa$ *is considered. It is not consistent with* $\{\neg\epsilon, \kappa \rightarrow \epsilon\}$*, so is not in the result.* $\epsilon$ *is similar. The set* $\{\kappa \rightarrow \epsilon\}$ *is a maximal subset of* $T$ *which is consistent with* $\neg\epsilon$. □

The previous definition describes how to obtain the epistemically consistent subset of a theory given an epistemic sequence. But there may be more than one epistemic sequence over a binary relation $\geq$. The following definition gives a resulting subtheory when there are multiple sequences.

**Definition 4.3** *Let* $T$ *be a partially closed theory and* $\mu$ *a set of sentences. Given a binary relation* $\geq$*, the* epistemical subtheory *of* $T$ *w.r.t.* $\mu$ *is defined as:*

$$ES(T, \mu) = \vee_{i=1}^{n} T_i$$

*where* $\{T_1, \ldots, T_n\}$ *is the set of all epistemically consistent subsets of* $T$ *w.r.t.* $\geq$. □

In other words, given multiple epistemic sequences, the epistemically consistent subset is derived from each one with respect to $\mu$, and the disjunction of all the subsets is taken. Obviously, $ES(T, \mu)$ is consistent with $\mu$.

At last we can define the update of a partially closed theory.

**Definition 4.4** *Let* $T$ *be a partially closed theory,* $\mu$ *be a set of sentences, and* $\geq$ *be a binary relation over* $T$*. Then the revision of* $T$ *by* $\mu$*, denoted as* $T\hat{+}\mu$*, is defined as*

$$(ES(T, \mu) \cup \mu)^{+}$$

□

The update of a partially closed theory $T$ is taken by defining a relation $\geq$ over all sentences in $T$. Using the priority relation a set of sequences is formed, and it is used to derive subtheories consistent with the new information. The disjunction of the subtheories is taken, and the new information is added to it. The partial closure of the result gives the result of the update.

An example is given to demonstrate the revision process.

**Example 4.5** *Let the sentence* $\gamma$ *represent the fact that it is cloudy outside,* $\delta$ *that it is raining,* $\zeta$ *that it is cold, and* $\varepsilon$ *that it is dark. The following hierarchy is used to show the relation* $\geq$*, where* $a \geq b$ *if the class of* $a$ *is less than that of* $b$.

*1.* $\gamma \rightarrow \delta \vee \zeta$

*2.* $\gamma \rightarrow \varepsilon, \gamma$

*3.* $\delta \vee \zeta$

*4.* $\delta, \zeta, \varepsilon$

*5.* $\neg\delta, \neg\zeta, \neg\varepsilon, \neg\gamma$

*The following assessment is our current knowledge:*

$$\{\gamma \rightarrow \delta \vee \zeta, \gamma \rightarrow \varepsilon\}$$

*Add the fact* $\gamma$ *to the knowledge. Since there is only one epistemic sequence the result is:*

$$\{\gamma \rightarrow \delta \vee \zeta, \gamma \rightarrow \varepsilon, \gamma, \delta \vee \zeta, \varepsilon\}$$

*At this point there are two epistemic sequences consistent with* $\geq$*:*

- $\gamma \rightarrow \delta \vee \zeta, \gamma \rightarrow \varepsilon, \gamma, \delta \vee \zeta, \varepsilon$
- $\gamma \rightarrow \delta \vee \zeta, \gamma, \gamma \rightarrow \varepsilon, \delta \vee \zeta, \varepsilon$

*Update the theory with* $\neg\varepsilon$*. The epistemically consistent subsets are:*

- $\{\gamma \rightarrow \delta \vee \zeta, \gamma \rightarrow \varepsilon, \delta \vee \zeta\}$
- $\{\gamma \rightarrow \delta \vee \zeta, \gamma, \delta \vee \zeta\}$

*The epistemic subtheory is* $\{\gamma \rightarrow \delta \vee \zeta, \delta \vee \zeta\}$ *and the result of the update is:*

$$\{\gamma \rightarrow \delta \vee \zeta, \delta \vee \zeta, \neg\varepsilon\}$$

□

It should be noted that the update sentences have implicit priority over all of the sentences in the theory when doing an update. The new sentences are never considered to be rejected from the theory when the update occurs. This is due to the second AGM postulate, which states that the new sentences should always exist after the update. It is easy to see how the update semantics could be extended if, for some reason, a need arises to consider the update sentences for derogation. The update sentences would be added to the original theory (without checks for consistency) and ordered over $\geq$ with all of the other sentences. Then the set would be considered for consistency. In fact, the proposed update semantics are equivalent to:

$$(\vee_{i=1}^{n} T_i')^+$$

where $T' = T \cup \mu$ and $\{T_1, \ldots, T_n\}$ is the set of all epistemically consistent subsets of $T'$ w.r.t. $\geq$ updated with $\emptyset$.

In this way it is possible to generalize the revision semantics to give "protected" formulas a higher priority than the update sentences, thus preserving rules which must always hold.

Many update operators based on a theory have suffered from problems of trying to find an appropriate solution when the update produces more than one result. By creating epistemical sequences over elements of a partially closed theory, then taking the disjunction of the results to be the base of the revision, we are guaranteed a meaningful result.

The revision is done on partially closed theories and it produces a desirable result. We hope that it can satisfy the revised AGM postulates, since they were designed to set the specifications for intuitive reasoning.

**Theorem 1** *The revision defined above satisfies* $R1$ - $R7$ *but does* not *satisfy* $R8$.

*Proof:* The revision satisfies $R1$ - $R7$.

$R1$: $T \hat{+} \mu$ is a partially closed theory.
Since $T^+$ is a partially closed theory for any set of sentences $T$ then $T \hat{+} \mu = (ES(T, \mu) \cup \mu)^+$ is a partially closed theory.

$R2$: $\mu \in T \hat{+} \mu$.
Since $T \hat{+} \mu = (ES(T, \mu) \cup \mu)^+$ then trivially $\mu \in T \hat{+} \mu$

$R3$: $T \cup \mu \models T \hat{+} \mu$.
Since we have that $\forall T_i' \subseteq T (T \models \vee_{i=1}^{n} T_i')$, then

$$T \models ES(T, \mu)$$
$$T \cup \mu \models (ES(T, \mu) \cup \mu)$$

$$T \cup \mu \models (ES(T,\mu) \cup \mu)^+$$
$$T \cup \mu \models T \hat{+} \mu$$

*R4*: $T \hat{+} \mu \models T \cup \mu$ if $T \cup \mu$ is consistent.
When $T \cup \mu$ is consistent then:

$$ES(T,\mu) = T$$
$$T \hat{+} \mu = (T \cup \mu)^+$$
$$T \hat{+} \mu \models T \cup \mu$$

*R5*: $T \hat{+} \mu$ consists of all propositional formulas only if $\mu$ is inconsistent.
Assume the contrary, that $T \hat{+} \mu$ consists of all propositional formulas and $\mu$ is consistent. When $\mu$ is consistent then $ES(T,\mu)$ is consistent also. But then $(ES(T,\mu) \cup \mu)^+$ is consistent. Contradiction.

*R6*: $(T \hat{+} \mu) = (T \hat{+} \alpha)$ if $\mu^+ = \alpha^+$.
We have that:

$$\mu^+ = \alpha^+$$
$$ES(T,\mu) = ES(T,\alpha)$$
$$(ES(T,\mu) \cup \mu^+) = (ES(T,\alpha) \cup \alpha^+)$$
$$(ES(T,\mu) \cup \mu)^+ = (ES(T,\alpha) \cup \alpha)^+$$

*R7*: $(T \hat{+} \mu) \cup \alpha \models T \hat{+} (\mu \cup \alpha)$.
Let $s_1, \ldots s_n$ be the $n$ epistemic sequences of $T$ given the relation $\geq$. Then there are $n$ epistemic subsets (Observation 1). We have that:

$$\forall s_i (sub(s_i, \mu \cup \alpha) \subseteq sub(s_i, \mu))$$
$$\vee_{i=1}^{n} sub(s_i, \mu) \models \vee_{i=1}^{n} sub(s_i, \mu \cup \alpha)$$
$$ES(T,\mu) \models ES(T, \mu \cup \alpha)$$
$$ES(T,\mu) \cup \mu \cup \alpha \models ES(T, \mu \cup \alpha) \cup \mu \cup \alpha$$
$$(ES(T,\mu) \cup \mu)^+ \cup \alpha \models (ES(T, \mu \cup \alpha) \cup \mu \cup \alpha)^+$$
$$(T \hat{+} \mu) \cup \alpha \models T \hat{+} (\mu \cup \alpha)$$

□

The next example demonstrates how postulate R8 is violated.

**Example 4.6** *Consider* $T = \{a, b \wedge c \rightarrow a, b, c, d \rightarrow c, d\}$, $\Lambda = \{\neg a\}$ *and* $\Omega = \{\neg d\}$ *with two epistemic sequences:*

$$b \wedge c \rightarrow a, d \rightarrow c, a, b, c, d$$

$$b \wedge c \rightarrow a, d \rightarrow c, c, d, a, b$$

*Then* $T\hat{+}\Lambda = \{b \wedge c \rightarrow a, d \rightarrow c, b \vee d, b \vee c, \neg b \vee \neg c, \neg a\}$. *Then* $T\hat{+}\Lambda$ *is consistent with* $\Omega$.

*Revise* $T$ *with* $\Lambda \cup \Omega$ *and the result is* $\{b \wedge c \rightarrow a, d \rightarrow c, b \vee c, \neg b \vee \neg c, \neg a, \neg d\}$. *Therefore,* $(T\hat{+}\Lambda) \cup \Omega \not\models T\hat{+}(\Lambda \cup \Omega)$

□

In [23] it was shown that *R8* may be too strong to be reasonable when applied to partially closed theories. Instead, the following condition was given:

$(R8')$ $(T\hat{+}\alpha) \wedge (T\hat{+}\mu) \models T\hat{+}(\alpha \vee \mu)$.

This postulate states that every world described by the knowledge base $T$ after revising it with $\Lambda$ and also with $\Omega$ must also be included in the worlds described by $T$ after revision with the disjunction of $\Lambda$ and $\Omega$.

It should be sufficient for the revision to satisfy *R8'* and not *R8*.

**Theorem 2** *The revision defined above satisfies R8'.*

*Proof:* We know that:

$$ES(T, \mu) \models ES(T, \mu \cup \alpha)$$

$$ES(T, \alpha) \models ES(T, \mu \cup \alpha)$$

$$ES(T, \mu) \cup ES(T, \alpha) \models ES(T, \mu \cup \alpha)$$

$$ES(T, \mu) \cup ES(T, \alpha) \cup \mu \cup \alpha \models ES(T, \mu \cup \alpha) \cup \mu \cup \alpha$$

$$(ES(T, \mu) \cup \mu)^{+} \cup (ES(T, \alpha) \cup \alpha)^{+} \models (ES(T, \mu \cup \alpha) \cup \mu \cup \alpha)^{+}$$

□

The revision operator has been shown to satisfy the revised Gārdenfors postulates, and also encompasses epistemic importance. If a reasonable epistemic importance can be placed on sentences in a theory, then the revision is straightforward and yields a meaningful result.

# Chapter 5

# The Application of Epistemic Importance

The last chapter gave an update algorithm which is applied to a partially closed theory with epistemic importance of facts already assigned. This section investigates how to assign epistemic importance to sentences such that reasonable and intuitive results are reached.

In [23] we gave an update semantics which satisfied the revised AGM postulates. The semantics assigned higher epistemic importance to rules than to arbitrary sentences. That work may be restated in terms the update semantics of last chapter with the following definition.

**Definition 5.1** *Let $T$ be a partially closed theory, and $\geq_{5.1}$ be the binary relation given by: $\forall \alpha \in T \, \forall \beta \in T (\alpha \geq_{5.1} \beta$ if there exists an atom $a$ such that $a$ is in $\beta$ and $\neg a$ is in $\alpha)$. An epistemic sequence may be defined in the usual way.* □

The binary relation on clauses is used to give higher priority to rules, since those should have highest priority to survive an update. An example from [23] will be used to demonstrate this.

**Example 5.1** *Let the database be $\{y \rightarrow s, y, s,\}$ and perform the following updates: $\{\neg s\}, \{g \vee b\}, \{\neg g\}, \{\neg b\}, \{\neg y\}$. The three possible epistemic sequences are initially:*

- $s, y \rightarrow s, y$

- $y \rightarrow s, s, y$
- $y \rightarrow s, y, s$

*and update by* $\{\neg s\}$ *yields* $\{y \rightarrow s, \neg s, \neg y\}$ *in all cases. The facts* $\{g \vee b\}$ *and* $\{\neg g\}$ *may be added to the database without conflict, so the database becomes* $\{y \rightarrow s, \neg y, \neg s, g \vee b, \neg g, b\}$. *All epistemic sequences have the ordering* $\neg g \geq_{5.1} g \vee b$, *so the revision with* $\{\neg b\}$ *yields* $\{y \rightarrow s, \neg y, \neg s, \neg b, \neg g\}$. *Adding* $\{\neg y\}$ *creates no change.* □

As well as assigning high priority to rules in a database, some other fundamentals should apply. The next definition gives an epistemic ordering to reflect such fundamentals.

**Definition 5.2** *Let* $T$ *be a partially closed theory, and* $\Lambda$ *update sentences. The notation* $IC(T)$ *is used to represent the set of integrity constraints in* $T$, $Lit(\alpha)$ *is used to represent the set of literals in clause* $\alpha$, *and* $NLit(\alpha) = \{\neg a \mid a \in Lit(\alpha)\}$. *Define the relation* $\geq_{5.2}$ *over* $T$ *as follows*

1. $\forall \gamma \in IC(T) \forall \alpha \notin IC(T)(\gamma \geq_{5.2} \alpha)$
2. $\forall \alpha \notin IC(T)(\Lambda \geq_{5.2} \alpha)$
3. $\forall \alpha \in \{(IC(T) \cup \Lambda)^{+} - IC(T)\} \forall \beta \notin \{IC(T) \cup \Lambda\}^{+}(\alpha \geq_{5.2} \beta)$
4. $\forall \alpha \in T - (IC(T) \cup \Lambda)^{+} \forall \beta \in T - (IC(T) \cup Lambda)^{+}(\alpha \geq_{5.1} \beta$ *if there exists an atom* $a$ *such that* $a$ *is in* $\beta$ *and* $\neg a$ *is in* $\alpha$*)*.
5. $\forall \alpha \in \{(T \dot{+} \Lambda) - T - (IC(T) \cup \Lambda)^{+}\} \forall \beta \in \{T \cap (T \dot{+} \Lambda)\}(\beta \geq_{5.2} \alpha)$

□

The first point is designed to give integrity constraints as least as much priority to survive derogation as anything else in the database. The second constraint places priority on the update information such that it has as least as much priority as anything that is not an integrity constraint. Next, priority is given to the minimally derivable logical consequences of the update sentences and integrity constraints. The fourth definition corresponds to Definition 5.1, but does not allow arbitrary rules to have priority over the integrity constraints or the new information, or minimally derivable consequences of the update sentences and integrity constraints. Condition five says that derived facts (excluding facts derived exclusively from the update sentences and integrity constraints) should have the lowest priority to survive derogation.

These ideas stem from a number of assumptions. One of these is that integrity constraints should have high priority. Another assumption is based upon the idea that one needs to keep track of the justifications of beliefs [11]. Although we do not keep track of justifications here, we place lower priority on beliefs which are derived from others, since they have less justification to survive after update. Another assumption is that inference rules should be more stable than basic facts during the update process.

The next example demonstrates an update semantics using priority relation $\geq_{5.2}$. This example will be used again in the next chapter when various update semantics are being compared.

**Example 5.2** *This problem is one of diagnostics. The knowledge base contains the integrity constraints: if one's stomach is sore (S) then they have the flu (F), or they have food poisoning (D), or they are pregnant (P). If a person is male (M), then they cannot be pregnant. If a person has not eaten lately (E), they cannot have food poisoning. So the knowledge base is:*

- $\{S \rightarrow F \vee D \vee P, M \rightarrow \neg P, \neg E \rightarrow \neg D\}$

*The first questioning reveals that a person has a sore stomach. The epistemic ordering after the update is (the epistemic order on the integrity constraints is not given, nor is it important for this example):*

- $S \rightarrow F \vee D \vee P, M \rightarrow \neg P, \neg E \rightarrow \neg D, S, F \vee D \vee P$

*The information $\neg F$ is then obtained:*

- $S \rightarrow F \vee D \vee P, M \rightarrow \neg P, \neg E \rightarrow \neg D, \neg F, S, F \vee D \vee P, D \vee P$

*Then $\neg E$ is determined:*

- $S \rightarrow F \vee D \vee P, M \rightarrow \neg P, \neg E \rightarrow \neg D, \neg E, \neg D, \neg F, S, F \vee D \vee P, D \vee P, P$

*The last piece of information obtained is $M$, and the result of the update is:*

- $\{S \rightarrow F \vee D \vee P, M \rightarrow \neg P, \neg E \rightarrow \neg D, M, \neg P, \neg E, \neg D, \neg F, \neg S\}$

□

In this example, $S$ is not in the result because it was given that $\neg F$ and $\neg D$ and $\neg P$, and $S$ was "older" than all of the other facts so has lowest priority to survive. This is due to the fact that we assume new knowledge is correct because it represents the latest information about the world. Therefore, the newer knowledge is, the more accurate we can assume that it is, hence it has higher priority over older knowledge.

The next example is 4.5 from the last chapter. In that example, the priorities of sentences were all fixed. This example will use the relation $\geq_{5.2}$ to define the priorities.

**Example 5.3** *Let the sentence $\gamma$ represent the fact that it is cloudy outside, $\delta$ that it is raining, $\zeta$ that it is cold, and $\varepsilon$ that it is dark. Given the knowledge:*

$$\{\gamma \rightarrow \delta \vee \zeta, \gamma \rightarrow \varepsilon\}$$

*And assume that $\gamma \rightarrow \delta \vee \zeta$ is the only integrity constraint. Add the fact $\gamma$ to the knowledge. The only resulting epistemic sequence is:*

$$\gamma \rightarrow \delta \vee \zeta, \gamma, \delta \vee \zeta, \gamma \rightarrow \varepsilon, \varepsilon$$

*Update the theory with $\neg\varepsilon$, and the epistemic sequence is:*

$$\gamma \rightarrow \delta \vee \zeta, \neg\varepsilon, \gamma, \delta \vee \zeta$$

□

Compare this with example 4.5, where after the first insertion there were two epistemic subtheories. Here there is only one because $\gamma \geq_{5.2} \delta \vee \zeta$ from 5.2[2] and $\delta \vee \zeta \geq_{5.2} \gamma \rightarrow \varepsilon$ from 5.2[3].

From the description of "update" and "revision" in [15], (see Chapter 3), we see that the priority given corresponds to the "revision" operation. Although the new update semantics best model a static world, they do place priority on newer facts and therefore incorporate more of a view of a changing world than a strict "revision" operator would.

# Chapter 6

# Comparisons of Update Semantics

This chapter is intended to give a feel for the various update semantics and to show the differences among them. The results are broken down into the model based semantics, the formula based semantics, and our own. It will be noted that the model based semantics and the formula based semantics both have strong points, depending upon the example used.

Our semantics is defined using the priority relation $\geq_{5.2}$ from the last chapter.

The following examples will be used, taken from various literature.

**Example 6.1** *[6]* $T = \{a \wedge b\}$, *and* $\Lambda = \{\neg a \vee \neg b\}$.
*The result should be* $\{(a \wedge \neg b) \vee (\neg a \wedge b)\}$. □

**Example 6.2** *[6]* $T = \{a \wedge b \wedge c\}$ *and* $\Lambda = \{(a \wedge \neg b \wedge \neg c) \vee (\neg a \wedge b \wedge c)\}$.
*The result of this update is not clear. Either it should be* $\Lambda$ *or it should be* $\{\neg a \wedge b \wedge c\}$. *The latter choice is more conservative and so represents the minimal change, but it is commonly accepted that all of the new information must be retained in an update.* □

Table 6.1 shows the results of examples 6.1 and 6.2 for the various update semantics.

| *Semantics* | *6.1* | *6.2* |
|---|---|---|
| Borgida | $(a \wedge \neg b) \vee (\neg a \wedge b)$ | $\neg a, b, c$ |
| Dalal | $(a \wedge \neg b) \vee (\neg a \wedge b)$ | $\neg a, b, c$ |
| Jackson/Pais | $(a \wedge \neg b) \vee (\neg a \wedge b)$ | $\Lambda$ |
| Satoh | $(a \wedge \neg b) \vee (\neg a \wedge b)$ | $\Lambda$ |
| Winslett | $(a \wedge \neg b) \vee (\neg a \wedge b)$ | $\Lambda$ |
| AGM | $\Lambda^*$ | $\Lambda^*$ |
| FUV | $\Lambda$ | $\Lambda$ |
| KUV | $\Lambda$ | $\Lambda$ |
| Ours | $\neg a \vee \neg b, a \vee b$ | $\Lambda$ |

Table 6.1: Comparisons of Update Semantics - Examples 6.1 and 6.2

**Example 6.3** *[14] $T = \{c \leftrightarrow a \vee b, c\}$ and $\Lambda = \{\neg c\}$, then $\Lambda = \{\neg a\}$.*

*Relate c to the fact that it is slippery, a that is it snowy, and b that it is frozen. After the update the conditions of it being slippery should not be lost. The final result should be $\{c \leftrightarrow a \vee b, \neg c, \neg a\}$.* □

**Example 6.4** *[21] This example is the well known blocks world. Consider a room with a TV and two air ducts. Let a denote the fact that the TV is on duct 1. Let b denote the fact that the TV is on duct 2, and c be that the TV is on the floor. The TV has to be on one of the spots, but can not be on more than one spot at one time. These facts can be represented by the database:*

$T = \{a \vee b \vee c, a \rightarrow \neg b, a \rightarrow \neg c, b \rightarrow \neg c\}$

*Update the database by the following:*

- $\{a \vee b\}$
- $\{a\}$
- $\{\neg a\}$

*In [21], it is claimed that all we can deduce after this is $\neg a$, and that b should not be assumed.* □

The results of examples 6.3 and 6.4 by all of the update semantics reviewed, as well as our own, are given in table 6.2.

| *Semantics* | *6.3* | *6.4* |
|---|---|---|
| Borgida | $\neg a, \neg c$ | $\neg a$ |
| Dalal | $\neg a, b, \neg c$ | $\neg a, \neg b, \neg c$ |
| Jackson/Pais | $\neg a, b, \neg c$ | $\neg a, (b \wedge \neg c) \vee (c \wedge \neg b)$ |
| Satoh | $\neg a, b, \neg c$ | $\neg a, \neg b, \neg c$ |
| Winslett | $\neg a, \neg c$ | $\neg a$ |
| AGM | $\{\neg c, \neg a\}^*$ | $\{\neg a\}^*$ |
| FUV | $c \leftrightarrow a \vee b, \neg a, \neg c$ | $T \cup \{a \vee b, \neg a\}$ |
| KUV | $c \leftrightarrow a \vee b, \neg a, \neg c$ | $T \cup \{a \vee b \leftrightarrow \neg a\}$ |
| Ours | $c \leftrightarrow a \vee b, \neg a, \neg b, \neg c$ | $IC(T) \cup \{\neg a, b \vee c, \neg b \vee \neg c\}$ |

Table 6.2: Comparisons of Update Semantics - Examples 6.3 and 6.4

A few of the examples will be described in detail for our update semantics.

**Example 6.5** *(example 6.2)*
*The theory and update sentence must be rephrased as a set of clauses, and so become:* $T = \{a, b, c\}$, $\Lambda = \{a \vee b, a \vee c, \neg b \vee \neg a, \neg b \vee c, \neg c \vee \neg a, \neg c \vee b\}$. *If the facts a, b, and c all have the same epistemic importance, then there are six epistemic sequences:*

- $\{a, b, c\}$
- $\{a, c, b\}$
- $\{b, c, a\}$
- $\{b, a, c\}$
- $\{c, a, b\}$
- $\{c, b, a\}$

*Then there are two epistemically consistent subsets of* $T$*, namely,* $\{a\}$*, and* $\{b, c\}$*. The result of the update is then* $\{a \vee b, a \vee c\} \cup \Lambda = \Lambda$ □

**Example 6.6** *( example 6.3 )*

*The theory is restated as a partially closed theory:* $\{\neg c \vee a \vee b, \neg a \vee c, \neg b \vee c, c, a \vee b\}$*, with the first three sentences considered as integrity constraints. Then an epistemic sequence is*

$$\neg c \vee a \vee b, \neg a \vee c, \neg b \vee c, c, a \vee b$$

*and the revision by* $\neg c$ *results in*

$$\{\neg c \vee a \vee b, \neg a \vee c, \neg b \vee c, \neg c, \neg a, \neg b\}$$

*which is also the result of the update with* $\neg a$*. Note that the facts* $\neg b$ *and* $\neg a$ *are logical consequences of the desired result (due to the partially closed theory).* □

**Example 6.7** *(example 6.4 ) There are four integrity constraints, which receive the highest epistemic importance. Initially, the four sentences may have any epistemic sequences. Choosing one:*

$$\neg a \vee \neg b, \neg a \vee \neg c, \neg b \vee \neg c, a \vee b \vee c$$

*After the update with* $\{a \vee b\}$*, the sequence is:*

$$\neg a \vee \neg b, \neg a \vee \neg c, \neg b \vee \neg c, a \vee b \vee c, a \vee b, \neg c$$

*Update with a and there are two sequences:*

$$\neg a \vee \neg b, \neg a \vee \neg c, \neg b \vee \neg c, a \vee b \vee c, a, \neg b, \neg c, a \vee b$$

$$\neg a \vee \neg b, \neg a \vee \neg c, \neg b \vee \neg c, a \vee b \vee c, a, \neg c, \neg b, a \vee b$$

*Then add* $\neg a$ *and there are two resulting epistemic sequences:*

$$\neg a \vee \neg b, \neg a \vee \neg c, \neg b \vee \neg c, a \vee b \vee c, \neg a, b \vee c, \neg b, c$$

$$\neg a \vee \neg b, \neg a \vee \neg c, \neg b \vee \neg c, a \vee b \vee c, \neg a, b \vee c, \neg c, b \vee c, b$$

*Which makes the result*

$$\{\neg a \vee \neg b, \neg a \vee \neg c, \neg b \vee \neg c, a \vee b \vee c, \neg a, b \vee c, \neg b \vee \neg c\}$$

□

Table 6.3 compares example 5.2 from the last chapter with the AGM and FUV approach, to demonstrate what impact the epistemic importance of sentences has on the outcome of our update, and also to show that the partially closed theory gives the correct result as compared to the closed theory (AGM) and a non-closed theory (FUV).

| *Semantics* | *5.2* |
|---|---|
| AGM | $\{M\}^*$ |
| FUV | $\{M\}$ |
| Ours | $\{S \rightarrow F \vee D \vee P, M \rightarrow \neg P, \neg E \rightarrow \neg D, M, \neg P, \neg E, \neg D, \neg F, \neg S\}$ |

Table 6.3: Update Comparisons - Theory Based

We are satisfied that all of the examples here yield a correct and intuitive result using our new update semantics and the given epistemic rules.

# Chapter 7

# Conclusion

There have been many update semantics proposed for deductive databases. Many of these semantics differ in the way that the database is viewed: some databases are viewed as a theory, and some are viewed as the underlying models. If a closed theory is considered then all derived facts are included, which may not be realistic. If the database is viewed in terms of the underlying models, then the meaning of the rules in the theory is lost. A problem with all of the approaches is how to deal with the epistemic importance of facts.

Clearly some sort of standard for deductive database updates is needed. Alchourron, Gärdenfors, and Makinson have done the first work in this area and propose criteria which must be met for every database update. This work has been instrumental in laying the foundations for updates. Since then update semantics have been proposed to meet the criteria. Most semantics have not been able to meet the AGM postulates, although some have been successful.

This thesis has proposed that the framework of the AGM postulates is not relevant to deductive database updates. Closed theories are unsuitable in the deductive database setting, because humans do not formulate all logical consequences of a set of beliefs. It is therefore unreasonable to do the same in a database. Instead, a new framework was proposed, that of partially closed theories. Partially closed theories contain a subset of the deductively closed set. Only those clauses which are minimally derived are added to the partially closed theory. This follows the way that a human would reason. This seems to solve the problems of the model theoretic approach and the formula theoretic approaches so far. In light of partially closed theories, the AGM postulates were modified.

As well, a new semantics was proposed. Given a binary relation on sentences in a theory, representing epistemic importance, the new semantics were defined, with partially closed theories as a basis.

The third and final contribution of this thesis was to give a set of rules to govern the assigning of epistemic importance to sentences of a theory.

The proof of any update algorithm is whether it satisfies the revised postulates, and how well it performs on various examples. The revision presented here was shown to satisfy the postulates. Numerous examples demonstrated the usefulness of the semantics.

It is hoped that the work done in this dissertation will help set the standards for deductive database updates.

# Bibliography


[1] S. Abiteboul. Updates a new frontier. In *International Conference on Database Theory*, pages 1–18, 1988.

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

[3] A. Borgida. Language features for flexible handling of exceptions in information systems. *ACM Transactions on Database Systems*, 10:565–603, 1985.

[4] C. Boutilier. A logic for revision and subjunctive queries. *Proc. of the 10th National (USA) Conf. on AI*, pages 609–615, 1992.

[5] M. Dalal. Investigations into a theory of knowledge base revision: Preliminary report. *7th National Conference on AI*, pages 1413–1419, 1988.

[6] M. Dalal. Updates in propositional databases. DCS tr 222, Department of Computer Science, Rutgers University, 1988.

[7] 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.

[8] W. A. Stubblefield G. F. Luger. *Artificial Intelligence and the Design of Expert Systems*. The Benjamin/Cummings Publishing Company, Inc., 1989.

[9] M. Y. Vardi G. M. Kuper, J. D. Ullman. On the equivalence of logical databases. *Proc. of the 3rd ACM SIGACT-SIGMOD Symposium on Principles of Database Systems*, pages 221–228, April 1984.

[10] P. Gärdenfors. Epistemic importance and minimal changes of belief. *Australasian Journal of Philosophy*, 62:137–157, 1984.

[11] P. Gärdenfors. *Knowledge in flux: modeling the dynamics of epistemic states.* The MIT Press, 1988.

[12] M. L. Ginsberg. Counterfactuals. *Artificial Intelligence*, 30:1:35 – 79, 1986.

[13] A. O. Mendelzon H. Katsuno. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52:263 – 294, 1989.

[14] A. O. Mendelzon H. Katsuno. A unified view of propositional knowledge base updates. *11th International Conference on AI*, pages 1413 – 1419, 1989.

[15] A. O. Mendelzon H. Katsuno. On the difference between updating a knowledge base and revising it. *Principles of Knowledge Representation and Reasoning: Proceedings of the 2nd International Conference*, pages 387 – 394, 1991.

[16] Hadley, Fagin, and Halpern. On logical omniscience: A critique with an alternative. In *Proceedings of the 6th Canadian Conf. on AI*, pages 49–5[illegible], 1986.

[17] J.W. Lloyd. *Foundations of Logic Programming (second edition).* Springer-Verlag, 1987.

[18] N. J. Nilsson M. R. Genesereth. *Logical Foundations of Artificial Intelligence.* Morgan Kaufmann Publishers, Inc., 1987.

[19] D. Makinson. How to give it up: A survey of some formal aspects of the logic of theory change. *Synthese*, 62:347–363, 1985.

[20] D. Makinson P. Gärdenfors. Revision of knowledge systems using epistemic entrenchment. *Proceeding of the 2nd International Conference on Theoretical Aspects of Reasoning*, pages 83 – 95, 1988.

[21] J. Pais P. Jackson. Semantic accounts of belief revision. In *Truth Maintenance Systems*, August 1990.

[22] K. Satoh. Nonmonotonic reasoning by minimal belief revision. *Proc. of the Intl. Conf. on 5th Generation Computer Systems*, pages 455–462, 1988.

[23] 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.

[24] M. Winslett. A framework for comparison of update semantics (extended abstract). *Proc. of the 7th ACM PODS*, pages 315–325, 1988.

[25] M. Winslett. Reasoning about action using a possible models approach. *Proc. of the 7th National (USA) Conf. on AI*, pages 89–93, 1988.

[26] M. Winslett. A model-theoretic approach to updating logical databases (extended abstract). *Proc. of the 5th ACM PODS*, pages 224–234, March 1986.