

A Unified View of Propositional Knowledge Base Updates

Hirofumi Katsuno*

NTT Basic Research Laboratories
Musashino-shi, Tokyo 180
Japan

Alberto O. Mendelzon†

Department of Computer Science
University of Toronto
Toronto, Ontario M5S 1A4
Canada

Abstract

The semantics of revising knowledge bases represented by sets of propositional sentences is analyzed from a model-theoretic point of view. A characterization of all revision schemes that satisfy the Gärdenfors rationality postulates is given in terms of an ordering among interpretations. Properties of the contraction operator that can be defined in terms of revision are also studied. Two new update operators, *elimination* and *recovery*, are introduced. Elimination discards all previous preconceptions on a set of propositional letters; recovery undoes the effect of the last update. It is shown that elimination cannot be expressed as a contraction, and that recovery is in general impossible. The existence of an invariant part of the knowledge base comprising a set of integrity constraints is considered and the definition of revision and contraction are modified to take integrity constraints into account.

1 Introduction

Consider a knowledge base (KB) represented by a set of sentences in a language L . As our perception of the world described by the knowledge base changes, the knowledge base must be revised. Several kinds of revisions may occur. If we simply acquire additional knowledge about the world, and the new knowledge does not conflict with the current beliefs¹ of the KB, there seems to be no difficulty — at least in theory — in incorporating the new knowledge in the form of new sentences. If, however, the new knowledge is inconsistent with the old beliefs, and we want the KB to be always consistent, we must resolve the conflict somehow; this operation will be called *revision*. A different kind of change occurs when a sentence previously believed becomes questionable; we call the operation that makes this change a *contraction*. A third operation erases all knowledge that involves a particular

fact; we call this *elimination*. Eliminating p will result not only in uncertainty over whether p is true or false, but we will also have to give up a belief that, say, $q \equiv p$. The fourth kind of change we consider is *retraction*: it involves undoing the effect of a previous operation.

Foundational work on knowledge base revision was done by Gärdenfors and his colleagues [Gärdenfors, 1984, Alchourrón *et al.*, 1985, Gärdenfors and Makinson, 1988]. They propose, on philosophical grounds, a set of *rationality postulates* that the operations of contraction and revision must satisfy and explore the implications of these postulates. The Gärdenfors postulates do not assume any concrete representation of the KB; in fact, KB's are modeled as deductively closed sets of sentences in some unspecified language. When we consider computer-based KB's, we need to fix a formalism and a finite syntactic representation of a KB. In this paper, we will assume the KB is represented by a finite set of propositional sentences. For this case, the both AI and database literatures contain several proposals on the appropriate definitions for some of the update operators [Dalal, 1988a, Dalal, 1988b, Fagin *et al.*, 1983, Weber, 1986, Winslett, 1987, Winslett, 1988, Borgida, 1985, Satoh, 1988]. The question now arises of whether the result of an update will depend on the particular set of sentences in the KB, or only on the worlds described by them. We are interested in methods that satisfy Dalal's *Principle of Irrelevance of Syntax*, that is, the meaning of the KB that results from an update must be independent of the syntax of the original KB, as well as independent of the syntax of the update itself.

Dalal [1988b] was the first one to relate his approach to the Gärdenfors postulates, pointing out that his proposal for the revision operator satisfies them. He did not analyze contraction, elimination, or retraction. In this paper, we go further by giving a model-theoretic characterization of all revision operators that satisfy the postulates. Our main theorem, in Section 3, shows that these operators are precisely the ones that accomplish an update with minimal change to the set of models of the KB. Dalal's method is seen to be a special case; we also discuss how the methods of Borgida, Winslett, Satoh, Weber, and Fagin, Ullman and Vardi fit into this framework.

Gärdenfors *et al.* show that the definition of a revision operator satisfying the rationality postulates uniquely

*This work was done while the first author was visiting the University of Toronto.

†This work was partially supported by the Natural Science and Engineering Research Council of Canada.

¹We use the terms *belief* and *knowledge* interchangeably in this paper.

determines a contraction operator that satisfies the corresponding postulates for contraction; conversely, revision can be defined in terms of contraction. In Section 4, we translate this definition to the case where the KB is a finite set of propositional sentences and show that contraction amounts to adding to the models of the KB the models of the KB revised with $\neg\mu$. We justify the definition further by showing that it is a sufficient condition for guaranteeing that old knowledge is not unnecessarily discarded.

It is generally recognized that not all sentences in a KB will have the same epistemic status. For example, an integrity constraint, or a definition of a concept in terms of others, should probably be treated differently than a fact about the domain. Gärdenfors and Makinson [1988] and Fagin, Ullman, and Vardi [1983] approach the problem in a similar way: rank the sentences in the KB according to their importance or “epistemic entrenchment” and take this into account when minimizing change from the old KB to the new one. In Section 5, we follow a different route. We distinguish between the KB and a set of integrity constraints IC and show how to modify the revision and contraction operators to ensure that the constraints are satisfied.

Finally, in Section 6 we treat a new operation, elimination, and show that it cannot in general be simulated by contraction, and in Section 7, we show that retraction is in general not achievable.

2 Preliminaries

Throughout this paper, we consider the language L of propositional logic, and we denote the set consisting of all the propositional letters in L by Ξ . We represent a knowledge base by a propositional formula ψ .

We use the standard terminology of propositional logic except for the definitions given below.

An interpretation of L is a function from Ξ to $\{T, F\}$. We often denote an interpretation by a tuple representing each propositional letter's value, e.g., if $\Xi = \{a, b, c, d\}$ then $\langle T, F, T, F \rangle$ is the interpretation which maps a, b, c, d to T, F, T, F respectively. A model of a propositional formula ψ is an interpretation that makes ψ true in the usual sense. $\text{Mod}(\psi)$ denotes the set of all the models of ψ .

Let $\mathcal{M}$ be a set of interpretations of L . Then, $\text{form}(\mathcal{M})$ denotes a formula whose set of models is equal to $\mathcal{M}$.

Let ψ be a propositional formula and let a be a propositional letter. Then, ψ_a^+ is defined as a formula obtained from ψ by replacing every occurrence of a with *true*, and ψ_a^- is defined as a formula obtained from ψ by replacing every occurrence of a with *false*. We define $\text{res}_a(\psi)$ as $\psi_a^+ \vee \psi_a^-$. If Ω is a set of propositional letters, $\{a_1, a_2, \dots, a_n\}$, $\text{res}_\Omega(\psi)$ is defined as $\text{res}_{a_1}(\text{res}_{a_2}(\dots(\text{res}_{a_n}(\psi))\dots))$.

Let I be an interpretation of L , and let Ω be a subset of Ξ . Then, $I|_\Omega$ is an interpretation over Ω obtained from I by restricting its domain to Ω . The complement of Ω in Ξ (i.e., $\Xi - \Omega$) is denoted by Ω^c .

3 Revision

Given a knowledge base ψ and a sentence μ , $\psi \circ \mu$ denotes the *revision* of ψ by μ ; that is, the new knowledge base obtained by adding new knowledge μ to the old knowledge base ψ .

3.1 The Gärdenfors Postulates for Revision

Gärdenfors and his colleagues propose the following postulates which they argue must be satisfied by any reasonable revision function. These postulates are formulated in a very general setting, but we restrict the discussion here to the propositional logic case. Instead of a finite KB, they consider a *knowledge set*, that is, a deductively closed set of formulas. Given knowledge set K and sentence μ , $K^*\mu$ is the revision of K by μ . $K^+\mu$ is the smallest deductively closed set containing K and μ . $K_\perp$ is the set consisting of all the propositional formulas.

- (G*1) $K^*\mu$ is a knowledge set.
- (G*2) $\mu \in K^*\mu$
- (G*3) $K^*\mu \subseteq K^+\mu$
- (G*4) If $\neg\mu \notin K$, then $K^+\mu \subseteq K^*\mu$
- (G*5) $K^*\mu = K_\perp$ only if μ is unsatisfiable.
- (G*6) If $\mu \equiv \phi$ then $K^*\mu = K^*\phi$.
- (G*7) $K^*(\mu \wedge \phi) \subseteq (K^*\mu)^+\phi$
- (G*8) If $\neg\phi \notin K^*\mu$ then $(K^*\mu)^+\phi \subseteq K^*(\mu \wedge \phi)$

If we fix a way of representing any knowledge set K by a propositional formula ψ such that $K = \{\phi \mid \psi \vdash \phi\}$, we can establish a direct correspondence between $K^*\mu$ and $\psi \circ \mu$. The following lemma characterizes revision operators that satisfy the first six postulates, (G*1)~(G*6).

Lemma 3.1 *Let $*$ be a revision operator on knowledge sets and $\circ$ its corresponding operator on KB's. Then $*$ satisfies (G*1)~(G*6) if and only if $\circ$ satisfies conditions (R1)~(R4) below.*

- (R1) $\psi \circ \mu$ implies μ .
- (R2) If $\psi \wedge \mu$ is satisfiable then $\psi \circ \mu \equiv \psi \wedge \mu$.
- (R3) If μ is satisfiable then $\psi \circ \mu$ is also satisfiable.
- (R4) If $\psi_1 \equiv \psi_2$ and $\mu_1 \equiv \mu_2$ then $\psi_1 \circ \mu_1 \equiv \psi_2 \circ \mu_2$.

This Lemma gives us a good grasp of the meaning of the first six postulates: new knowledge (μ) is retained in the updated KB (R1), his Principle of Irrelevance of Syntax (R4), a guarantee that the obvious path will be taken when there is no conflict (R2), and a condition preventing a revision from introducing unwarranted inconsistency (R3). What about the remaining two postulates? The following lemma rephrases them in terms of $\circ$.

Lemma 3.2 *(G*7) and (G*8) are equivalent to (R5) and (R6) respectively in the same sense as Lemma 3.1.*

- (R5) $(\psi \circ \mu) \wedge \phi$ implies $\psi \circ (\mu \wedge \phi)$.
- (R6) If $(\psi \circ \mu) \wedge \phi$ is satisfiable then $\psi \circ (\mu \wedge \phi)$ implies $(\psi \circ \mu) \wedge \phi$.

To grasp the intuitive meaning of (R5) and (R6), consider the set of models of the KB, $Mod(\psi)$. Suppose that there is some metric for measuring the “distance” between $Mod(\psi)$ and any interpretation I . We want our revision operator to effect minimal change, that is, we want the models of $\psi \circ \mu$ to be those models of μ which are closest to $Mod(\psi)$ with respect to our distance metric.

Rule (R5) says that our notion of closeness is well-behaved in the sense that if we pick any interpretation I which is closest to $Mod(\psi)$ in a certain set, namely $Mod(\mu)$, and I also belongs to a smaller set, $Mod(\mu \wedge \phi)$, then I must also be closest to $Mod(\psi)$ within the smaller set $Mod(\mu \wedge \phi)$.

A violation of rule (R6) would imply that an interpretation I may be closer to the KB than J within a certain set, while J is closer than I within some other set. To see this, consider a model I of $\psi \circ (\mu \wedge \phi)$, that is, a model of $\mu \wedge \phi$ that is closest to $Mod(\psi)$. Suppose I is not a model of $(\psi \circ \mu) \wedge \phi$. The precondition of (R6) says that there is some interpretation J that is a model of $\psi \circ \mu$ and also of ϕ . That is, J is a model of ϕ that is closest to $Mod(\psi)$. But then J is closer to $Mod(\psi)$ within the set $Mod(\phi)$ than I , while I is closer to $Mod(\psi)$ than J within the set $Mod(\mu \wedge \phi)$. In the next section we formalize the notion of an interpretation being closer to the KB than another one and relate it to the postulates.

3.2 Orders between Interpretations

The postulates (G*7) and (G*8) represent the condition that revision be accomplished with minimal change. In this Subsection, we give a model theoretic characterization of minimal change.

Let $\mathcal{I}$ be the set of all the interpretations of L . A *pre-order* $\leq$ over $\mathcal{I}$ is a reflexive and transitive relation on $\mathcal{I}$. Consider a function that assigns to each propositional formula ψ a pre-order $\leq_\psi$ over $\mathcal{I}$. We say this assignment is persistent if the following three conditions hold:

1. If $I \in Mod(\psi)$ then for all $I' \in \mathcal{I}$, $I \leq_\psi I'$ holds.
2. If $I \in Mod(\psi)$ and $I' \notin Mod(\psi)$ then $I' \leq_\psi I$ does not hold.
3. If $\psi \equiv \phi$, then $\leq_\psi = \leq_\phi$.

That is, every model of ψ is less than or equal to every other model and no non-model can be less than or equal to a model. We define $<_\psi$ as $I <_\psi I'$ if and only if $I \leq_\psi I'$ and $I' \not\leq_\psi I$.

Let $\mathcal{M}$ be a subset of $\mathcal{I}$. An interpretation I is minimal in $\mathcal{M}$ with respect to $\leq_\psi$ if $I \in \mathcal{M}$ and there is no $I' \in \mathcal{M}$ such that $I' <_\psi I$. Let

$$Min(\mathcal{M}, \leq_\psi) = \left\{ I \mid \begin{array}{l} I \text{ is minimal in } \mathcal{M} \\ \text{with respect to } \leq_\psi \end{array} \right\}.$$

If we regard $\leq_\psi$ as a measure representing the closeness between $Mod(\psi)$ and an interpretation, i.e., $I' \leq_\psi I$ means that I' is closer to $Mod(\psi)$ than I , then $Min(\mathcal{M}, \leq_\psi)$ can be seen as the set of all the closest interpretations in $\mathcal{M}$ to $Mod(\psi)$.

Theorem 3.1 *Revision operator $\circ$ satisfies Conditions (R1)~(R6) if and only if there exists a persistent assignment that maps each KB ψ to a total pre-order $\leq_\psi$ such that $Mod(\psi \circ \mu) = Min(Mod(\mu), \leq_\psi)$.*

The persistent assignment of the Theorem maps KB's to total pre-orders. If we allow two interpretations to be incomparable under some pre-order, then existence of the pre-order is no longer sufficient to guarantee condition (R6). Instead of (R6), two weaker conditions characterize the existence of a persistent assignment to partial pre-orders. A later version of this paper will provide the details.

3.3 Review of Proposals from the Literature

3.3.1 Dalal's Revision

Dalal [1988a, 1988b] uses the number of propositional letters on which two interpretations differ as a measure of “distance” between them. This distance measure induces an ordering among interpretations as follows. First, define the distance between two interpretations I and J , $dist(I, J)$, as the total number of propositional letters whose interpretation is different in I and J . Next, define the distance between $Mod(\psi)$ and I as

$$dist(Mod(\psi), I) = \min_{J \in Mod(\psi)} dist(J, I).$$

Then, we can define a persistent assignment of a total pre-order $\leq_\psi$ as $I \leq_\psi J$ if and only if

$$dist(Mod(\psi), I) \leq dist(Mod(\psi), J).$$

And Dalal's revision operator $\circ_D$ can be defined by:

$$Mod(\psi \circ_D \mu) = Min(Mod(\mu), \leq_\psi).$$

Thus, it follows from Theorem 3.1 that Dalal's revision operator $\circ_D$ satisfies Conditions (R1)~(R6).

Example 3.1 Let L have only four propositional letters, a, b, c, d . Consider the following five interpretations:

$$\begin{aligned} I_1 &= \langle T, T, T, T \rangle, & I_2 &= \langle F, F, F, F \rangle, \\ J_1 &= \langle F, F, T, T \rangle, & J_2 &= \langle T, F, F, F \rangle, \\ J_3 &= \langle F, F, T, F \rangle. \end{aligned}$$

Let

$$\begin{aligned} \psi &\equiv form(I_1, I_2), & \phi_1 &\equiv form(J_1, J_2, J_3), \\ \phi_2 &\equiv form(J_1, J_2), & \phi_3 &\equiv form(J_1, J_3). \end{aligned}$$

Then, we can obtain

$$\begin{aligned} \psi \circ_D \phi_1 &\equiv form(J_2, J_3), & \psi \circ_D \phi_2 &\equiv form(J_2), \\ \psi \circ_D \phi_3 &\equiv form(J_3). \end{aligned}$$

3.3.2 Borgida's Revision

Let $\circ_B$ be the revision operator proposed by Borgida [1985] and extended in [Dalal, 1988a]. He concentrates on sets of propositional letters on which a model of ψ and a model of μ differ.

We say that two interpretations, I and J , differ on a set of propositional letters, Ω , if Ω is the set of propositional letters a such that the truth value of a in I is

different from its truth value in J . $\text{Diff}(I, \mu)$ is the collection of all the sets of propositional letters on which I and some model of μ differ. Then, Borgida's revision operator $\circ_B$ is defined as follows. If μ is inconsistent with ψ , then an interpretation J is a model of $\psi \circ_B \mu$ if and only if J is a model of μ and there is some model I of ψ such that the set of propositional letters on which I and J differ is a minimal element of $\text{Diff}(I, \mu)$. Otherwise, i.e., if μ is consistent with ψ , then $\psi \circ \mu$ is defined as $\psi \wedge \mu$.

If ψ has only one model I , we can represent the models of $\psi \circ_B \mu$ as the set of minimal elements of the partial order $\leq_I$ defined by $J_1 \leq_I J_2$ if and only if the set on which I and J_1 differ is a subset of the set on which I and J_2 differ. This fact makes us expect that $\circ_B$ might be defined in terms of a persistent assignment of a partial pre-order to each KB. However, the following example shows that $\circ_B$ cannot be defined in this way.

Example 3.2 Consider Example 3.1 again. We add the following two interpretations:

$$J_4 = \langle T, T, F, F \rangle, \quad J_5 = \langle T, T, T, F \rangle.$$

Let

$$\begin{aligned} \phi_4 &\equiv \text{form}(J_2, J_4), & \phi_5 &\equiv \text{form}(J_4, J_5), \\ \phi_6 &\equiv \text{form}(J_2, J_4, J_5). \end{aligned}$$

Then, we can obtain

$$\begin{aligned} \psi \circ_B \phi_4 &\equiv \text{form}(J_2, J_4), & \psi \circ_B \phi_5 &\equiv \text{form}(J_4, J_5), \\ \psi \circ_B \phi_6 &\equiv \text{form}(J_2, J_5). \end{aligned}$$

Suppose that there is some persistent assignment of a partial pre-order $\leq_\psi$ that captures the $\circ$ operator. Then, $J_2 \not\leq_\psi J_4$ follows from $\psi \circ_B \phi_4 \equiv \text{form}(J_2, J_4)$. $J_5 \not\leq_\psi J_4$ does from $\psi \circ_B \phi_5 \equiv \text{form}(J_4, J_5)$. On the other hand, either $J_2 <_\psi J_4$ or $J_5 <_\psi J_4$ follows from $\psi \circ_B \phi_6 \equiv \text{form}(J_2, J_5)$. This is a contradiction.

In terms of Conditions (R1)~(R6), the following Lemma holds

Lemma 3.3 Assume that ψ is consistent. Then, Borgida's revision operator $\circ_B$ satisfies (R1)~(R5), but does not satisfy (R6).

3.3.3 Winslett's Revision

Winslett [1988] proposes a revision operator which is suitable for reasoning about action. Her revision operator is defined for the first order calculus case and called the *possible models* approach. We restrict this operator to the propositional case and denote it by $\circ_{pma}$.

If the new knowledge μ is inconsistent with the old knowledge base ψ , then Winslett's operator $\circ_{pma}$ coincides with Borgida's operator $\circ_B$. However, even if μ is consistent with ψ , Winslett defines $\circ_{pma}$ in the same way as the inconsistent case. This means that $\circ_{pma}$ violates Condition (R2), i.e., when $\psi \wedge \mu$ is consistent, $\psi \circ_{pma} \mu$ may not be equivalent to $\psi \wedge \mu$.

Example 3.3 Let L have only two propositional letters, a and b . Let $\psi \equiv (a \wedge \neg b) \vee (\neg a \wedge b)$ and $\mu \equiv a$. Then, we obtain $\psi \circ_{pma} \mu \equiv a$ and $\psi \circ_B \mu \equiv (a \wedge \neg b)$.

To interpret this example in the context of [Winslett, 1988], suppose a represents the fact that a book is on the floor and b means that a magazine is on the floor. Then, ψ states that either the book or the magazine is on the floor, but not both. Now, we order a robot to put the book on the floor. The result of this action should be represented by the revision of ψ with a . After the robot puts the book on the floor, we know a , but we do not know whether b is true or not. Hence, the new state of the world should be described not as $a \wedge \neg b$ (Borgida's revision), but as a (Winslett's revision).

In terms of Conditions (R1)~(R6), the following holds.

Lemma 3.4 Assume that ψ is consistent. Then, Winslett's revision operator $\circ_{pma}$ satisfies (R1) and (R3)~(R5), but violates (R2) and (R6).

3.3.4 Satoh's Revision

Satoh [1988] proposed a revision operator for first order knowledge bases by using the notion of circumscription. If we apply his revision operator to the propositional logic case, we obtain the revision operator $\circ_S$ which corresponds to a global version of Borgida's revision operator.

Satoh considers the minimal elements of $\text{Diff}(\psi, \mu)$, where

$$\text{Diff}(\psi, \mu) = \cup_{I \in \text{Mod}(\psi)} \text{Diff}(I, \mu).$$

He defines an interpretation J to be a model of $\psi \circ_S \mu$ if and only if J is a model of μ and there is some model I of ψ such that the set on which I and J differ is a minimal set of $\text{Diff}(\psi, \mu)$.

This definition makes the closeness between models of ψ and models of μ depend on both ψ and μ . Hence, we cannot expect $\circ_S$ to be defined in terms of a persistent assignment of a pre-order to each KB. The following example shows that this expectation is correct.

Example 3.4 Consider Example 3.1 again. We obtain $\psi \circ_S \phi_1 \equiv \text{form}(J_2, J_3)$, $\psi \circ_S \phi_2 \equiv \phi_2$, and $\psi \circ_S \phi_3 \equiv \phi_3$.

Suppose there is some persistent assignment of a pre-order $\leq_\psi$ that captures the $\circ_S$ operator. Then, $\psi \circ_S \phi_1 \equiv \text{form}(J_2, J_3)$ implies either $J_2 <_\psi J_1$ or $J_3 <_\psi J_1$. However, $J_2 \not\leq_\psi J_1$ follows from $\psi \circ \phi_2 \equiv \phi_2$, and $J_3 \not\leq_\psi J_1$ follows from $\psi \circ_S \phi_3 \equiv \phi_3$, a contradiction.

In terms of Conditions (R1)~(R6), the following holds.

Lemma 3.5 Assume that ψ is consistent. Then, Satoh's revision operator $\circ_S$ satisfies (R1)~(R5), but does not satisfy (R6).

3.3.5 Weber's Revision

Weber [1986] also concentrates on sets of propositional letters on which a model of ψ and a model of μ differ. His revision operator $\circ_W$ can be easily defined by elimination, which we discuss in Section 6. Thus, we only point out that $\circ_W$ cannot be defined in terms of a persistent assignment of a pre-order to each KB, since $\circ_D$ yields the same results as $\circ_S$ in the case of Example 3.4.

In terms of Conditions (R1)~(R6), the following holds.

Lemma 3.6 Assume that both ψ and μ consistent. Then, Weber's revision operator $\circ_W$ satisfies (R1)~(R5), but does not satisfy (R6).

3.3.6 Fagin, Ullman and Vardi's Revision

The approach of Fagin, Ullman and Vardi [1983], when applied to deductively closed sets of sentences, yields a revision operator $\circ_F$ such that, if ψ and μ are inconsistent, $\psi \circ_F \mu$ is equivalent to μ . Hence, $\circ_F$ can be defined by the persistent assignment that maps each ψ to the pre-order $\leq_\psi$ such that $I \leq_\psi I'$ if and only if either I is a model of ψ or neither I nor I' is a model of ψ .

Because it is not very satisfactory to throw away all the old knowledge each time an inconsistent update is attempted, these authors favour sets of sentences that are not deductively closed; however, in this case their method produces a result which depends on the syntax of the KB, violating the Principle of Irrelevance of Syntax.

4 Contraction

When a sentence μ that was previously believed becomes uncertain, we apply a contraction operator to the KB to ensure that μ is not implied by the updated KB. Note that this is different from revising the KB with $\neg\mu$. Given a KB ψ and a sentence μ , $\psi \ominus \mu$ denotes the new knowledge base obtained by contracting μ from ψ . Gärdenfors and his colleagues have analyzed contraction carefully [Alchourrón *et al.*, 1985, Gärdenfors and Makinson, 1988], but the database and AI literature has concentrated much more on revision.

4.1 Connection between Contraction and Revision

Alchourrón, Gärdenfors, and Makinson [1985] also propose rationality postulates for contraction, (G-1)~(G-8). According to their notation, $K^- \mu$ is the new knowledge set obtained from the old knowledge set K by contracting μ .

(G-1) $K^- \mu$ is a knowledge set.

(G-2) $K^- \mu \subseteq K$

(G-3) If $\mu \notin K$, then $K^- \mu = K$.

(G-4) If μ is not a tautology, then $\mu \notin K^- \mu$.

(G-5) $K \subseteq (K^- \mu)^* \mu$

(G-6) If $\mu \equiv \phi$, then $K^- \mu = K^- \phi$

(G-7) $(K^- \mu \cap K^- \phi) \subseteq K^- (\mu \wedge \phi)$

(G-8) If $\mu \notin K^- (\mu \wedge \phi)$, then $K^- (\mu \wedge \phi) \subseteq K^- \mu$

We refer to [Makinson, 1985] for the justification of these postulates. If we represent each knowledge set by a propositional formula, as we did in Section 3.1, the following properties easily follow from (G-1)~(G-6).

(C-1) If ψ implies μ and μ is not a tautology, then $\psi \ominus \mu$ does not imply μ and ψ implies $\psi \ominus \mu$; otherwise, $\psi \ominus \mu$ is equivalent to ψ .

(C-2) $(\psi \ominus \mu) \wedge \mu$ implies ψ .

(C-1) shows that $\ominus$ realizes contraction if μ is implied by ψ and that otherwise $\ominus$ does not influence the KB.

(C-2) shows that the contraction of μ from ψ is strong enough to recover, when conjoined with μ , all the facts in the original KB ψ .

Alchourrón *et al.* [1985] showed that contraction and revision are closely related and can in fact be defined in terms of each other. They proved that, given a revision operator $*$ that satisfies (G*1)~(G*8), if we define the contraction operator $-$ by

$$K^- \mu = K \cap K^* \neg \mu, \quad (R \rightarrow C)$$

then contraction operation $-$ satisfies (G-1)~(G-8). Conversely, given a contraction operator $-$ that satisfies (G-1)~(G-8), if we define revision $*$ by

$$K^* \mu = (K^- \neg \mu)^+ \mu, \quad (C \rightarrow R)$$

then this revision operator satisfies (G*1)~(G*8).

If we rewrite $(R \rightarrow C)$ in our terminology, we get $\psi \ominus \mu \equiv (\psi \vee \psi \circ \neg \mu)$, because the set consisting of both ψ and $\psi \circ \neg \mu$ is inconsistent when ψ implies μ . This means that the result of contracting μ can be obtained by adding some models of $\neg \mu$ to the models of the old knowledge base. Adding some models of $\neg \mu$ is necessary to guarantee that $\psi \ominus \mu$ does not imply μ . However, we might think it wise to also add some models of $\neg \psi \wedge \mu$ or delete some models of ψ . The following Theorem shows this is not desirable; the requirement that the new KB be obtained by adding only models of $\neg \mu$ to the models of ψ is necessary and sufficient for guaranteeing that old knowledge ϕ is recoverable from $\psi \ominus \mu$ together with μ .

Theorem 4.1 The following conditions are equivalent.

1. ψ implies ϕ if and only if $\psi \ominus \mu$ implies $\mu \supset \phi$.
2. $Mod(\psi) = Mod(\psi \ominus \mu) \cap Mod(\mu)$.

4.2 Complex Contraction

It is not hard to imagine using compound sentences in a revision, e.g. "it will rain or snow tomorrow." It might be harder, at first glance, to imagine a need to contract a complex sentence. However, consider the following example.

Example 4.1 Let a represent the fact that roads are covered with snow, let b represent the fact that roads are frozen. Suppose that we know roads are slippery if and only if they are covered with snow or frozen. Suppose we know there have been no accidents all day, casting doubt on our belief that roads are slippery. The way to update the knowledge base is precisely to contract $a \vee b$ (i.e., $\psi \ominus (a \vee b)$).

Note 4.1 Example 4.1 brings out an interesting point. Let c represent the fact that roads are slippery. Suppose that we have the knowledge base ψ defined by $c \wedge (c \equiv a \vee b)$. That is, we know that roads are slippery if and only if they are covered with snow or ice. If we use the contraction function $\ominus$ defined from Dalal's revision function, $\psi \ominus c$ is equivalent to $a \vee b$. Hence, by contracting the knowledge that the roads are slippery, we also lose the notion of what "slippery" means. This example suggests that we are not quite modelling things in the right way. The next section shows how to do it right.

5 Integrity Constraints

Knowledge bases contain not only sentences of fact about the world, but also *integrity constraints* that are intended to ensure that the KB is an appropriate representation of the world. In the deductive database literature, it is common to regard constraints as sentences expressed in the same language as the KB. Reiter [1988] argues that it is more appropriate to represent constraints in an epistemic modal logic when the KB is a set of first-order sentences. We adopt a different point of view here; we express knowledge and constraints in the same language, but treat them differently under update. Integrity constraints are considered invariant, and updates are restricted to produce KB's which are consistent with the constraints. Winslett [1988] and Satoh [1988] take a similar approach. We present a modification of revision and contraction that accomplishes this and avoids the unintuitive results of Note 4.1.

Let IC be a propositional formula which represents an update-invariant component of the knowledge base. That is, we want the KB to imply IC at all times. (Note this is not the only possible way to model integrity constraints; see [Reiter, 1988] for a discussion of other ways.) For each revision function $\circ$ and each contraction function $\ominus$, we define $\circ^{IC}$ and $\ominus^{IC}$ by

$$\psi \circ^{IC} \mu \equiv \psi \circ (\mu \wedge IC)$$

and

$$\psi \ominus^{IC} \mu \equiv \psi \ominus (IC \supset \mu).$$

Intuitively, when revising the KB with μ we make sure that the constraints will hold in the result by attaching IC to μ . When contracting μ , we avoid an accidental contraction of IC by replacing μ with $IC \supset \mu$.

The following theorem shows that $\circ^{IC}$ and $\ominus^{IC}$ inherit desirable properties about $\circ$ and $\ominus$, respectively. We consider the following condition corresponding to (R3).

(R3') if $\mu \wedge IC$ is satisfiable then $\psi \circ^{IC} \mu$ is also satisfiable.

Theorem 5.1

1. If $\circ$ satisfies conditions (R1)~(R6), then $\circ^{IC}$ does (R1), (R2), (R3'), (R4)~(R6). In particular, if ψ implies IC then $\psi \circ^{IC} \mu$ implies both μ and IC .
2. Similarly, $\ominus^{IC}$ inherits properties of $\ominus$. In particular, if $\ominus$ satisfies the conditions corresponding to (G-1)~(G-6), then
 - (a) if ψ implies μ and IC does not imply μ then $\psi \ominus^{IC} \mu$ does not imply μ .
 - (b) if ψ does not imply μ then $\psi \ominus^{IC} \mu$ is equivalent to ψ .
 - (c) if ψ implies IC then $\psi \ominus^{IC} \mu$ also implies IC .

Before this section, we have abided by what Dalal calls the *Principle of Primacy of Update*: after revising KB ψ with sentence μ , the new KB must imply μ . This assumption is almost universal in the literature. In practice, it is of course unreasonable to expect the KB to blindly assimilate any new fact without questioning it. It is much more likely that a KB manager will seriously object to incorporating a sentence that violates

an integrity constraint. Now that we consider integrity constraints explicitly, we see that primacy of update no longer holds; if μ is inconsistent with IC , then $\circ^{IC}$ is undefined (or produces the inconsistent KB). Similarly, if IC implies μ , it is impossible to contract μ by using $\ominus^{IC}$.

Example 5.1 We use the KB ψ from Note 4.1. Let IC be $c \equiv (a \vee b)$. Suppose that both $\circ^{IC}$ and $\ominus^{IC}$ are defined by Dalal's revision function $\circ_D$. Then, $\psi \circ^{IC} \neg c$ (i.e., $\psi \circ_D (IC \wedge \neg c)$) is equivalent to $(c \equiv a \vee b) \wedge \neg a \wedge \neg b \wedge \neg c$, and $\psi \ominus^{IC} c$ (i.e., $\psi \vee \psi \circ_D \neg (IC \supset c)$) is equivalent to $c \equiv a \vee b$.

6 Elimination

Elimination is used to discard all the knowledge concerning Ω where Ω is a set of propositional letters. We denote the new knowledge base obtained by elimination as $\psi \oslash \Omega$. Then, the new knowledge base should have the following properties.

(E1) Let Γ be the set of all propositional formulas, ϕ , such that ϕ is implied by ψ and that no letter in Ω occurs in ϕ . Let σ be a formula in which some letter in Ω occurs. Then, Γ implies σ if and only if $\psi \oslash \Omega$ implies σ .

(E2) Let ϕ be a formula in which no letter in Ω occurs. Then, ψ implies ϕ if and only if $\psi \oslash \Omega$ implies ϕ .

Intuitively, (E1) means that $\psi \oslash \Omega$ loses all the non-trivial information about Ω . (E2) means that $\psi \oslash \Omega$ does not lose any information about Ω^c .

The following theorem shows that Conditions (E1) and (E2) uniquely determine the elimination function.

Theorem 6.1 $\psi \oslash \Omega$ satisfies Conditions (E1) and (E2) if and only if $\psi \oslash \Omega \equiv \text{res}_{\Omega}(\psi)$.

Weber's revision operator $\circ_W$ we referred to in Section 3.3.5 is now defined as follows. If neither ψ nor μ is unsatisfiable, then $\psi \circ_W \mu$ is defined by $(\psi \oslash \Omega) \wedge \mu$, where Ω is the union of the minimal sets of $\text{Diff}(\psi, \mu)$. If either ψ or μ is unsatisfiable, then $\psi \circ_W \mu$ is defined as ψ .

Winslett [1987] proposes another revision operator $\circ_M$ defined by $\psi \circ_M \mu \equiv (\psi \oslash \Omega) \wedge \mu$ where Ω is the set consisting of all the propositional letters which occur in μ . The revision operator violates the Principle of Irrelevance of Syntax because Ω depends on the syntactic representation of μ . She tries to represent contraction by revising the knowledge base with a tautology, e.g., $\psi \ominus a$ is represented by $\psi \circ_M (a \vee \neg a)$.

The following theorem shows that elimination cannot be represented by contraction in general.

Theorem 6.2 Let $\circ$ be a update operator defined by a persistent assignment of total pre-order $\leq_{\psi}$ to each KB ψ . Let $\ominus$ be the contraction operator defined by $\psi \ominus \mu \equiv (\psi \vee \psi \circ \neg \mu)$ corresponding to the equation $(R \rightarrow C)$. Let Ω be a set of propositional letters. Then, the following two conditions are equivalent.

1. There is no μ such that $\psi \ominus \mu \equiv \psi \oslash \Omega$.
2. There are two interpretations, I and I' , such that

- (a) $I <_{\psi} I'$,
- (b) I is not a model of ψ , but there is some model of ψ , J , satisfying $J|_{\Omega^c} = I|_{\Omega^c}$,
- (c) $I|_{\Omega^c} = I'|_{\Omega^c}$

Example 6.1 Let ψ is equivalent to $a \wedge (b \supset c) \wedge (c \supset d) \wedge (d \supset b)$. Let $\Omega = \{a, b\}$. Then, $\psi \odot \Omega$ is equivalent to $c \equiv d$.

On the other hand, let $\ominus$ be the contraction function defined by Dalal's revision function. Let $I = \langle F, F, F, F \rangle$ and $I' = \langle F, T, F, F \rangle$. Then, I and I' satisfy Condition (2) of Theorem 6.2. Hence, for any μ , $\psi \ominus \mu$ is not equivalent to $\psi \odot \Omega$.

7 Recovery

Suppose we change our mind about an update we just performed and want to take it back; this is what we call the recovery operation. How can it be implemented?

First, consider the recovery of revision. Let $\circ$ be a revision function that retains complete new knowledge. Suppose that ψ is any formula and μ is a formula that has exactly one model. Then, it is easy to show $\psi \circ \mu \equiv \mu$. In this case, the new knowledge base depends only on the added new knowledge and is independent of the old knowledge base. Therefore, we cannot in general recover the old knowledge base ψ from the added new knowledge and the updated knowledge base.

On the other hand, if a contraction function satisfies (G-1)~(G-6) then the following holds.

1. If ψ implies μ then $\psi \equiv (\psi \ominus \mu) \wedge \mu$.
2. If ψ does not imply μ then $\psi \equiv \psi \ominus \mu$. Note that $\psi \not\equiv (\psi \ominus \mu) \wedge \mu$.

Hence, we can recover the previous contraction only if we know whether μ was implied by the old knowledge or not.

8 Conclusion

Our main result is a model-theoretic characterization of revision schemes that satisfy Gärdenfors' rationality postulates in the propositional case. Future research directions include extending these results to the first order predicate calculus and applying the results about revision to areas such as diagnosis from first principles, abductive reasoning, and database updates.

We have also presented some basic properties of contraction and introduced the operations of elimination and recovery. We have shown that recovery is in general unachievable; a future research direction is how to carry along syntactic information in the knowledge base in order to make recovery possible.

Acknowledgements

We would like to thank Marianne Winslett for clarifying the difference between her approach and Satoh's, and Mukesh Dalal for pointing out our earlier misunderstanding of Borgida's approach. We are grateful to Ray Reiter for introducing us to this area.

References

- [Alchourrón et al., 1985] Carlos E. Alchourrón, Peter Gärdenfors and David Makinson. On the logic of theory change: partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [Borgida, 1985] Alex Borgida. Language Features for Flexible Handling of exceptions in information systems. *ACM Trans. on Database Systems*, 10:563–603, 1985.
- [Dalal, 1988a] Mukesh Dalal. Updates in propositional databases. Technical Report DCS-TR-222, Department of Computer Science, Rutgers University, Feb. 1988.
- [Dalal, 1988b] Mukesh Dalal. Investigations into a theory of knowledge base Revision: Preliminary Report. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, pages 475–479, 1988.
- [Fagin et al., 1983] Ronald Fagin, Jeffrey D. Ullman and Moshe Y. Vardi. On the semantics of updates in databases. In *Proceedings of the Second ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 352–365, 1983.
- [Gärdenfors, 1984] Peter Gärdenfors. Epistemic importance and minimal changes of belief. *Australasian Journal of Philosophy*, 62:136–157, 1984.
- [Gärdenfors and Makinson, 1988] Peter Gärdenfors and David Makinson. Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 83–95, 1988.
- [Makinson, 1985] David Makinson. How to give it up: A survey of some formal aspects of the logic of theory change. *Synthese*, 62:347–363, 1985.
- [Reiter, 1988] Raymond Reiter. On integrity constraints. in *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 97–111, 1988.
- [Satoh, 1988] Ken Satoh. Nonmonotonic reasoning by minimal belief revision. In *Proceedings of the International Conference on Fifth Generation Computer Systems*, pages 455–462, 1988.
- [Weber, 1986] Andreas Weber. Updating propositional formulas. In *Proceedings of the First Conference on Expert Database Systems*, pages 487–500, 1986.
- [Winslett, 1987] Marianne Winslett. Updating databases with incomplete information. Technical Report STAN-CS87-1143, Stanford University, 1987.
- [Winslett, 1988] Marianne Winslett. Reasoning about action using a possible models approach. In *Proceedings of the Seventh National Conference on Artificial Intelligence*, pages 89–93, 1988.