A Consistency-Based Model for Belief Change: Preliminary Report

James P. Delgrande
School of Computing Science,
Simon Fraser University
Burnaby, B.C., Canada V5A 1S6
jim@cs.sfu.ca

Torsten Schaub
Institut für Informatik,
Universität Potsdam
D–14415 Potsdam, Germany
torsten@cs.uni-potsdam.de

Abstract

We present a general, consistency-based framework for belief change. Informally, in revising K by α, we begin with α and incorporate as much of K as consistently possible. Formally, a knowledge base K and sentence α are expressed, via renaming propositions in K, in separate alphabets, but such that there is an isomorphism between the original and new alphabets. Using a maximization process, we assume that corresponding atoms in each language are equivalent insofar as is consistently possible. Lastly, we express the resultant knowledge base using just the original alphabet. There may be more than one way in which α can be so extended by K: in *choice revision*, one such "extension" represents the revised state; alternately *revision* consists of the intersection of all such extensions.

The overall framework is flexible enough to express other approaches to revision and update, and the incorporation of static and dynamic integrity constraints. Our framework differs from work based on ordinal conditional functions, notably with respect to iterated revision. We argue that the approach is well-suited for implementation: choice revision gives better complexity results than general revision; the approach can be expressed in terms of a finite knowledge base; and the scope of a revision can be restricted to just those propositions mentioned in the sentence for revision α.

Introduction

We describe a general framework for belief change. The approach has something of the same flavour as the consistency-based paradigm for diagnosis (Reiter 1987) or the assumption-based approach to default reasoning (Poole 1988), although it differs significantly in details. Informally, in revising a knowledge base K by sentence α, we begin with α and incorporate as much of K as consistently possible. There may be more than one way in which information from K can be incorporated. This gives rise to two notions of revision: a choice notion, in which one such "extension" is used for the revised state, and the intersection of all such extensions. Belief contraction is defined analogously.

We mainly focus on belief revision in this paper. For revision, first a knowledge base K and sentence α are expressed, via renaming atomic propositions in K, in separate alphabets. We next assume that as many atoms in α are equivalent to the corresponding atom in K, as consistently possible. A set of such equivalent atoms is used to incorporate as much of the original knowledge base as is consistently possible. In the final section we discuss the more general approach, which we show is flexible enough to express extant approaches to revision and update.

The approach is developed in a formal, abstract framework. However, we argue that it is well-suited for implementation: The notion of choice revision gives better complexity results than general revision; moreover, we argue that belief revision is an area in which choice reasoning makes sense in some cases. Second, we show how the approach can be expressed equivalently in terms of a finite knowledge base, in place of a deductively-closed belief set. Third, we show that the scope of a revision can be restricted to just those propositions common to the knowledge base and sentence for revision.

We begin by presenting a very general framework for expressing belief change. This is restricted to address revision and contraction. Following this, we show how the approach allows for a uniform treatment of integrity constraints. As well, the approach supports iterated revision, with properties distinct from approaches based on the work of Spohn (Spohn 1988). Finally we briefly explore the general framework, and suggest it is flexible enough to express extant approaches to revision and update.

Background

A common approach in belief revision is to provide a set of *rationality postulates* for revision and contraction functions. The *AGM approach* of Alchourron, Gärdenfors, and Makinson (Gärdenfors 1988), provides the best-known set of such postulates. The goal is to describe belief change on an abstract level, independent of how beliefs are represented and manipulated. Belief states, called *belief sets*, are modelled by sets of sentences closed under the logical consequence operator of some logic in some language $\mathcal{L}$, where the logic includes classical propositional logic. For belief set K, $K+\alpha$ is the deductive closure of $K \cup \{\alpha\}$, and is called the *expansion* of K by α. $K_\perp$ is the inconsistent belief set (i.e. $K_\perp = \mathcal{L}$). $\mathcal{T}$ is the set of all belief sets.

A *revision* function $\dot{+}$ is a function from $\mathcal{T} \times \mathcal{L}$ to $\mathcal{T}$ satisfying the following postulates.

($K\dot{+}1$) $K\dot{+}\alpha$ is a belief set.

($K\dot{+}2$) $\alpha \in K\dot{+}\alpha$.

($K\dot{+}3$) $K\dot{+}\alpha \subseteq K + \alpha$.

($K\dot{+}4$) If $\neg\alpha \notin K$, then $K + \alpha \subseteq K\dot{+}\alpha$.

($K\dot{+}5$) $K\dot{+}\alpha = K_\perp$ iff $\vdash \neg\alpha$.

($K\dot{+}6$) If $\vdash \alpha \equiv \beta$, then $K\dot{+}\alpha = K\dot{+}\beta$.

($K\dot{+}7$) $K\dot{+}(\alpha \wedge \beta) \subseteq (K\dot{+}\alpha) + \beta$.

($K\dot{+}8$) If $\neg\beta \notin K\dot{+}\alpha$, then $(K\dot{+}\alpha) + \beta \subseteq K\dot{+}(\alpha \wedge \beta)$.

That is: the result of revising K by α is a belief set in which α is believed; whenever the result is consistent, revision consists of the expansion of K by α; the only time that $K_\perp$ is obtained is when α is inconsistent; and revision is independent of the syntactic form of K and α. The last two postulates deal with the relation between revising with a conjunction and expansion.

(Katsuno & Mendelzon 1992) explores the distinct notion of belief *update* in which an agent changes its beliefs in response to changes in its external environment. Our interests here centre on revision; however as the end of the paper, we briefly consider this approach.

Recently there has been interest in *iterated* belief revision, a topic that the AGM approach by-and-large leaves open. Representative work includes (Boutilier 1994; Williams 1994; Lehmann 1995; Darwiche & Pearl 1997). We discuss Darwiche and Pearl's approach here. They employ the notion of an *epistemic state* that encodes how the revision function changes following a revision. Ψ denotes an epistemic state; $Bel(\Psi)$ denotes the belief set corresponding to Ψ. So now the result of revising an epistemic state is another epistemic state (from which the revised belief set may be determined using $Bel(\cdot)$. Darwiche and Pearl propose the following postulates that "*any* rational system of belief change should comply with" (p. 2). Following their practice, we use Ψ to stand for $Bel(\Psi)$ when it appears as an argument of $\models$.

$C1$: If $\alpha \models \mu$ then $(\Psi\dot{+}\mu)\dot{+}\alpha \equiv \Psi\dot{+}\alpha$.

$C2$: If $\alpha \models \neg\mu$ then $(\Psi\dot{+}\mu)\dot{+}\alpha \equiv \Psi\dot{+}\alpha$.

$C3$: If $\Psi\dot{+}\alpha \models \mu$ then $(\Psi\dot{+}\mu)\dot{+}\alpha \models \mu$.

$C4$: If $\Psi\dot{+}\alpha \not\models \neg\mu$ then $(\Psi\dot{+}\mu)\dot{+}\alpha \not\models \neg\mu$.

(Nayak *et al.* 1996) propose a variant of $C2$ along with the following postulate:

$Conj$: If $\alpha \wedge \beta \not\models \perp$ then $(\Psi\dot{+}\alpha)\dot{+}^{\alpha}\beta = \Psi\dot{+}(\alpha \wedge \beta)$.

where $\dot{+}^{\alpha}$ indicates that the change in $\dot{+}$ following revision by α depends in part on α. This postulate is shown to be strong enough to derive $C1$, $C3$, and $C4$ in the presence of the other postulates.

There has also been work on specific approaches to revision based on the distance between models of a knowledge base and a sentence to be incorporated in the knowledge base. This work includes (Dalal 1988; Forbus 1989; Satoh 1988; Winslett 1988). In these approaches, models of the new knowledge base consist of models of the sentence to be added that are closest (based on "distance" between atomic sentences) to models of the original knowledge base.

Our approach differs from previous work first, in that we provide a specific, albeit general, framework in which approaches may be expressed. As well, the general framework allows the incorporation of different forms of integrity constraints. Also, given that it falls into the "consistency-based" paradigm, the approach has a certain syntactic flavour. However, notably, our approach is independent of the syntactic form of the knowledge base and sentence for revision.

Our technique of maximizing sets of equivalences of propositional letters bears a superficial resemblance to the use of such equivalences in (Liberatore & Schaerf 1997) (based in turn on the technique developed in (de Kleer & Konolige 1989)). However the approaches are distinct; in particular and in contradistinction to these references, we employ disjoint alphabets for a knowledge base and revising sentence. As well, the approach bears a resemblance to that of (del Val 1993). However, unlike del Val, we provide a single approach which may be restricted to yield extant approaches; also, we place no a priori restrictions on the form of a knowledge base.

Formal Preliminaries

We deal with propositional languages and use the logical symbols $\top$, $\perp$, $\neg$, $\vee$, $\wedge$, $\supset$, and $\equiv$ to construct formulas in the standard way. We write $\mathcal{L}_{\mathcal{P}}$ to denote a language over an alphabet $\mathcal{P}$ of *propositional letters* or *atomic propositions*. Formulas are denoted by the Greek letters α, β, α_1, ... *Knowledge bases* or, equivalently, *belief sets* are initially identified with deductively-closed sets of formulas and are denoted K, K_1, ... So we have $K = Cn(K)$, where $Cn(\cdot)$ is the deductive closure of the formula or set of formulas given as argument. Later we relax this restriction.

Given an alphabet $\mathcal{P}$, we define a disjoint alphabet $\mathcal{P}'$ as $\mathcal{P}' = \{p' \mid p \in \mathcal{P}\}$. Then, for $\alpha \in \mathcal{L}_{\mathcal{P}}$, we define α' as the result of replacing in α each proposition p from $\mathcal{P}$ by the corresponding proposition p' in $\mathcal{P}'$ (so implicitly there is an isomorphism between $\mathcal{P}$ and $\mathcal{P}'$). This is defined analogously for sets of formulas.

We define a *belief change scenario* in language $\mathcal{L}_{\mathcal{P}}$ as a triple $B = (K, U, V)$, where K, U, V are sets of formulas in $\mathcal{L}_{\mathcal{P}}$. Informally, K is a knowledge base that will be changed such that the set U will be true in the result, and the set V will be consistent with the result. For a base approach to revision we take $V = \emptyset$ and for a base approach to contraction we take $U = \emptyset$.

In the definition below, "maximal" is with respect to set containment (rather than set cardinality). The following is our central definition.

Definition 1 *Let $B = (K, U, V)$ be a belief change scenario in $\mathcal{L}_{\mathcal{P}}$. Define EQ as a maximal set of equivalences $EQ \subseteq \{p \equiv p' \mid p \in \mathcal{P}\}$ such that*

$$K' \cup EQ \cup U \cup V \not\vdash \perp.$$

Then

$$Cn(K' \cup EQ \cup U) \cap \mathcal{L}_{\mathcal{P}}$$

is a consistent definitional extension *of B.*

Hence, a consistent definitional extension of B is a modification of K in which U is true, and in which V is consistent.

We say that EQ *underlies* the consistent definitional extension of B. We let $\overline{EQ}$ stand for $\{p \equiv p' \mid p \in \mathcal{P}\} \setminus EQ$.

Clearly, for a given belief change scenario there may be more than one consistent definitional extension. We will make use of the notion of a *selection function* c that for any set $I \neq \emptyset$ has as value some element of I. In Definition 2 and 3, these primitive functions can be regarded as inducing selection functions c' on belief change scenarios, such that $c'((K, U, V))$ has as value some consistent definitional extension of (K, U, V). This is a slight generalisation of selection functions as found in the AGM approach.

Revision and Contraction

Definition 1 provides a very general framework for specifying belief change. In the next two definitions we give specific definitions for revision and contraction. We develop these specific approaches and then, at the end of the paper, we return to the more general framework of Definition 1 and discuss how it can be used to express other approaches.

Definition 2 (Revision) *Let K be a knowledge base and α a formula, and let $(E_i)_{i\in I}$ be the family of all consistent definitional extensions of $(K, \{\alpha\}, \emptyset)$. Then*

1. $K \dot{+}_c \alpha = E_i$ *is a* choice revision *of K by α with respect to some selection function c with $c(I) = i$.*
2. $K \dot{+} \alpha = \bigcap_{i\in I} E_i$ *is the* (skeptical) revision *of K by α.*

Table 1 gives examples of (skeptical) revision. The first column gives the original knowledge base, but with atoms already renamed. The second column gives the revision formula, while the third gives the EQ set(s) and the last column gives the results of the revision. For the first and last column, we give a formula whose deductive closure gives the corresponding belief set.

K'	α	EQ	$K \dot{+} \alpha$
$p' \wedge q'$	$\neg q$	$\{p \equiv p'\}$	$p \wedge \neg q$
$\neg p' \equiv q'$	$\neg q$	$\{\, p \equiv p',\ q \equiv q'\,\}$	$p \wedge \neg q$
$p' \vee q'$	$\neg p \vee \neg q$	$\{\, p \equiv p',\ q \equiv q'\,\}$	$p \equiv \neg q$
$p' \wedge q'$	$\neg p \vee \neg q$	$\{p \equiv p'\},\ \{q \equiv q'\}$	$p \equiv \neg q$

Table 1: (Skeptical) revision examples.

In detail, for the last example, we wish to determine

$$\{p \wedge q\} \dot{+} (\neg p \vee \neg q)\ . \tag{1}$$

We find maximal sets $EQ \subseteq \{p \equiv p', q \equiv q'\}$ such that

$$\{p' \wedge q'\} \cup EQ \cup \{\neg p \vee \neg q\} \cup \emptyset \text{ is consistent.}$$

We get two such sets of equivalences, namely $EQ_1 = \{p \equiv p'\}$ and $EQ_2 = \{q \equiv q'\}$. Accordingly, we obtain

$$\{p \wedge q\} \dot{+} (\neg p \vee \neg q) = \bigcap_{i=1,2} Cn(\{p' \wedge q'\} \cup EQ_i \cup \{\neg p \vee \neg q\}) \cap \mathcal{L}_{\mathcal{P}}.$$

In addition to $(\neg p \vee \neg q)$, we get $(p \vee q)$, jointly implying $(p \equiv \neg q)$.

In this example we get two choice extensions, $Cn(p \wedge \neg q)$ and $Cn(\neg p \wedge q)$. This raises the question of the usefulness of choice revision compared to general revision. An apparent limitation of a choice reasoner is that it might draw overly strong conclusions. However, in belief revision this may be less of a problem than, say, in non-monotonic reasoning: the goal in revision is to determine the true state of the world; if a (choice) revision results in an inaccurate knowledge base, then *this* inaccuracy will presumably be detected and rectified in a later revision. So choice revision may do no worse than a "skeptical" operator with respect to "converging" to the true state of the world. In addition, as we later show, it may do so significantly more efficiently and with better worst-case behaviour. Hence for a land vehicle exploring a benign environment, choice revision might be an effective part of a control mechanism; for something like flight control, or controlling a nuclear reactor, one would prefer skeptical revision.

Contraction is defined similarly to revision.

Definition 3 (Contraction) *Let K be a knowledge base and α a formula, and let $(E_i)_{i\in I}$ be the family of all consistent definitional extensions of $(K, \emptyset, \{\neg\alpha\})$. Then*

1. $K \dot{-}_c \alpha = E_i$ *is a* choice contraction *of K by α with respect to some selection function c with $c(I) = i$.*
2. $K \dot{-} \alpha = \bigcap_{i\in I} E_i$ *is the* (skeptical) contraction *of K by α.*

Table 2 gives examples of (skeptical) contraction, using the same format and conventions as Table 1.

K'	α	EQ	$K \dot{-} \alpha$
$p' \wedge q'$	q	$\{p \equiv p'\}$	p
$p' \wedge q' \wedge r'$	$p \vee q$	$\{r \equiv r'\}$	r
$p' \vee q'$	$p \wedge q$	$\{\, p \equiv p',\ q \equiv q'\,\}$	$p \vee q$
$p' \wedge q'$	$p \wedge q$	$\{p \equiv p'\},\ \{q \equiv q'\}$	$p \vee q$

Table 2: (Skeptical) contraction examples.

In detail, for the first example we wish to determine

$$\{p \wedge q\} \dot{-} q\ . \tag{2}$$

We compute the consistent definitional extensions of $(\{p \wedge q\}, \emptyset, \{\neg q\})$. We rename the propositions in $\{p \wedge q\}$ and look for maximal subsets EQ of $\{p \equiv p', q \equiv q'\}$ such that

$$\{p' \wedge q'\} \cup EQ \cup \emptyset \cup \{\neg q\} \text{ is consistent.}$$

We obtain $EQ = \{p \equiv p'\}$, yielding

$$\begin{aligned}\{p \wedge q\} \dot{-} q &= Cn(\{p' \wedge q'\} \cup \{p \equiv p'\} \cup \emptyset) \cap \mathcal{L}_{\mathcal{P}} \\ &= Cn(\{p\}).\end{aligned}$$

Properties of Revision and Contraction

With respect to the AGM postulates, we obtain the following.

Theorem 1 *Let $\dot{+}$ and $\dot{+}_c$ be defined as in Definition 2. Then $\dot{+}$ and $\dot{+}_c$ satisfy the basic AGM postulates $(K\dot{+}1)$ to $(K\dot{+}4)$, $(K\dot{+}6)$ as well as $(K\dot{+}7)$.*

For $(K\dot{+}5)$ we have instead the weaker postulate:

$(K\dot{+}5)$ $K \dot{+} \alpha = K_\perp$ iff: $K = K_\perp$ or $\vdash \neg\alpha$.

We obtain analogous results for $\dot{-}$ and $\dot{-}_c$ with respect to the AGM contraction postulates:

Theorem 2 *Let $\dot{-}$ and $\dot{-}_c$ be defined as in Definition 3. Then $\dot{-}$ satisfies the basic AGM postulates $(K\dot{-}1)$ to $(K\dot{-}4)$, $(K\dot{-}6)$, and $(K\dot{-}7)$. In addition, $\dot{-}_c$ satisfies the basic AGM postulates $(K\dot{-}1)$ to $(K\dot{-}4)$, $(K\dot{-}6)$.*

We also obtain the following interdefinability results:

Theorem 3 (Levi Identity) $K\dot{+}\alpha = (K\dot{-}\neg\alpha) + \alpha$.

Theorem 4 (Partial Harper Identity)
$K\dot{-}\alpha \subseteq K \cap (K\dot{+}\neg\alpha)$

The following example shows that equality fails in the Harper Identity: if $K \equiv p \wedge q \wedge r$ and $\alpha \equiv r$, then $K\dot{-}\alpha \equiv r$ while $K \cap (K\dot{+}\neg\alpha) \equiv (p \equiv \neg q) \wedge r$. Similar results are obtained for choice revision and choice contraction by appeal to appropriate selection functions.

Iterated belief change: The approach obviously supports iterated revision. Since we use a "global" metric, and since we can assume that *every* revision result, given K and α, can be determined, we continue to use K here rather than Darwiche and Pearl's Ψ for an epistemic state. That is, for us, we don't need to refer to epistemic states, since we have completely specified how $\dot{+}$ should behave on all arguments. Nonetheless, neither operator in Definition 2 satisfies any of the Darwiche-Pearl postulates for iterated revision. Nor in our opinion should they. For example, for $C1$, if we have

$$K = Cn(\neg p), \quad \alpha = p, \quad \mu = p \vee q, \tag{3}$$

then in our approach we obtain that

$$(K\dot{+}\mu)\dot{+}\alpha = Cn(p \wedge q) \text{ but } K\dot{+}\alpha = Cn(p). \tag{4}$$

(Darwiche & Pearl 1997; Nayak *et al.* 1996) assert that these results should be equal. However, it is *possible* (contra $C1$) that there are cases where revising $\neg p$ by $p \vee q$ yields $\neg p \wedge q$ and a subsequent revision by p then gives $p \wedge q$, but revising $\neg p$ by p would yield p. Which is to say, a significant difficulty in the area of belief revision is that different people have conflicting intuitions. However, Darwiche and Pearl argue that *all rational* revision functions should obey $C1$. Consequently they would need to argue that in all cases, having (4) result from (3) is *irrational*.

More seriously, an instance of $C2$ (letting α be $\neg\phi$ and μ be $\phi \wedge \psi$, whence $\alpha \models \neg\mu$) is the following:

$C2'$**:** $(K\dot{+}(\phi \wedge \psi))\dot{+}\neg\phi \equiv K\dot{+}\neg\phi$.

Thus if you revise by $(\phi \wedge \psi)$ and then revise by the negation of some of this information ($\neg\phi$), then the other original information (ψ) is lost. So, in a variant of an example from (Darwiche & Pearl 1997), consider where I see a new bird in the distance and come to believe that it is red and flies. If on closer examination I see that it is yellow, then according to $C2'$ and so $C2$, I also no longer believe that it flies. This seems too strong a condition to want to adopt. We conjecture (but have no proof) that approaches based on (Spohn 1988), such as (Darwiche & Pearl 1997), are subject in some form to such a "blanketing" result.

On the other hand, there are nontrivial results concerning iterated revision that hold for the present approach. For example, we have:

Theorem 5 *Let $\dot{+}$ be defined as in Definition 2. Then:* $(\alpha\dot{+}\beta)\dot{+}\alpha = \beta\dot{+}\alpha$.

Semantics: The operator $\dot{+}$ provides a (near) syntactic counterpart to the minimal-distance-between-models approach of (Satoh 1988). For two sets S and T, let $S\Delta T$ be the symmetric difference, $(S \cup T) \setminus (S \cap T)$. For formulas α, β, define

$$\Delta^{\min}(\alpha, \beta) = \min_{\subseteq}(\{M\Delta M' \mid M \in Mod(\alpha), M' \in Mod(\beta)\}) \,,$$

where $Mod(\alpha)$ is the set of all models of α, each of which is identified with a set of propositions. Then, we have:

Theorem 6 *Let $B = (K, U, \emptyset)$ be a belief change scenario in $\mathcal{L}_{\mathcal{P}}$ where $K \neq \mathcal{L}_{\mathcal{P}}$, and let $(EQ_i)_{i \in I}$ be the family of all sets of equivalences, as defined in Definition 1.*

Then, $\{\ \{p \in \mathcal{P} \mid (p \equiv p') \notin EQ_i\} \mid i \in I\} = \Delta^{\min}(U, K)$.

This correspondence is interesting, but is of limited use beyond supplying a semantics for one instance of the approach. The choice approach, and (below) considerations on implementation and integrity constraints, are not readily expressed in Satoh's model-based semantics. As well, a contraction function is straightforwardly obtained in Satoh's approach only by using the Harper Identity (which doesn't fully obtain here). Further, in the last section, we show how other approaches can be expressed in our general framework.

Integrity Constraints

Definitions 2 and 3 are similar in form, differing only in how the formula α is mapped onto the sets U and V in Definition 1. Clearly one can combine these definitions, allowing simultaneous revision by one formula and contraction by another. This in-and-of-itself isn't overly interesting, but it does lead to a natural and general treatment of integrity constraints in our approach.

There are two standard definitions of a knowledge base K satisfying a static integrity constraint IC. In the *consistency-based* approach of (Kowalski 1978), K satisfies IC iff $K \cup \{IC\}$ is satisfiable. In the *entailment-based* approach of (Reiter 1984), K satisfies IC iff $K \vdash IC$. (Katsuno & Mendelzon 1991) show how entailment-based constraints can be maintained across revisions: given an integrity constraint IC and revision function $\dot{+}$, a revision function $\dot{+}^{IC}$ which preserves IC is defined by: $K\dot{+}^{IC}\alpha = K\dot{+}(\alpha \wedge IC)$. In our approach, we can define revision taking into account both approaches to integrity constraints.

Corresponding to Definition 2 (and ignoring the choice approach) we obtain:

Definition 4 *Let K be a knowledge base, α a formula, and IC_K, IC_R sets of formulas. Let $(E_i)_{i \in I}$ be the family of all consistent definitional extensions of $(K, \{\alpha\} \cup IC_R, IC_K)$. Then $K\dot{+}^{(IC_K, IC_R)}\alpha = \bigcap_{i \in I} E_i$ is the* revision *of K by α incorporating integrity constraints IC_K (consistency-based) and IC_R (entailment-based).*

Theorem 7 *Let $\dot{+}^{(IC_K, IC_R)}$ be defined as in Definition 4. Then $\left(K\dot{+}^{(IC_K, IC_R)}\alpha\right) \vdash IC_R$. If $IC_R \cup IC_K \not\vdash \neg\alpha$ then $\left(K\dot{+}^{(IC_K, IC_R)}\alpha\right) \cup IC_K$ is satisfiable.*

Finally, and in contrast with previous approaches, it is straightforward to add *dynamic* integrity constraints, which express constraints that hold between states of the knowledge base before and after revision. The simplest way of so doing is to add such constraints to the set V in Definition 1. To state that if $a \wedge b$ is true in a knowledge base before revision then c must be true afterwards, we would add $a' \wedge b' \supset c$ to V. Note however that the addition of dynamic constraints may lead to an operator that violates some of the properties of $\dot{+}$. For example $Cn(\alpha)\dot{+}\neg\alpha$ with dynamic constraint $\alpha' \supset \alpha$ leads to an inconsistent revision.

Implementability Considerations

We claimed at the outset that the approach is well-suited for implementation. To this end, we first consider the use of choice belief revision. Second we consider the problem of representing the results of revision in a finite, manageable representation. Lastly, we address limiting the range of EQ.

Complexity: From (Eiter & Gottlob 1992) and Theorem 6 it follows that deciding, for given K, α, β, whether $K\dot{+}\alpha \vdash \beta$ is Π_2^P-complete. However, the analogous problem for choice revision is lower in the polynomial hierarchy.

Theorem 8 *Given a selection function c, formulas K, α, β, and a set of equivalnces EQ. Then, we have:*

1. *Deciding whether EQ determines a choice revision of K and α is in $\mathsf{P}^{\mathbf{NP}}$.*
2. *Deciding $K\dot{+}_c\alpha \vdash \beta$ is in $\mathsf{P}^{\mathbf{NP}}$.*

We have not yet addressed restrictions on the syntactic form of K or α; but see (Eiter & Gottlob 1992).

Finite representations: Definitions 1, 2, and 3 provide a characterisation of revision and contraction, yielding in either case a deductively-closed belief set. Here we consider how the same (with respect to logical equivalence) operators can be defined, but where a knowledge base is given as an arbitrary, finite set of formulas. It follows from the discussion below that, for knowledge base K and formula α, we can defined choice revision so that $|K\dot{+}_c\alpha| \leq |K| + |\alpha|$ for any selection function c.

Informally the procedure is straightforward, although the technical details are less so. A knowledge base K is now represented by a formula (or set of formulas). Via Definitions 1 and 2 we consider maximal sets EQ where $\{K'\} \cup \{\alpha\} \cup EQ$ is consistent. For each such set EQ, we replace each p' in K' by p where $(p \equiv p') \in EQ$ and we replace each p' in K' by $\neg p$ where $(p \equiv p') \in \overline{EQ}$. The result of these substitutions into $\{K'\} \cup \{\alpha\}$ is a sentence of size $\leq |K| + |\alpha|$ and whose deductive closure is equivalent to (some) choice revision. The disjunction of all such sentences (and so considering all possible sets EQ) is equivalent to $Cn(K)\dot{+}\alpha$.

As opposed to the computation of the sets EQ, the result of revising or contracting a formula K can be captured without an explicit change of alphabet. We start by observing that any set of equivalences EQ induces a binary partition of its underlying alphabet $\mathcal{P}$, namely $\langle \mathcal{P}_{EQ}, \mathcal{P}_{\overline{EQ}}\rangle$ with $\mathcal{P}_{EQ} = \{p \in \mathcal{P} \mid p \equiv p' \in EQ\}$ and $\mathcal{P}_{\overline{EQ}} = \mathcal{P} \setminus \mathcal{P}_{EQ}$. Given a belief change scenario B along with a set of equivalences EQ_i (according to Definition 1), we define for $\alpha \in \mathcal{L}_{\mathcal{P}}$, that $\lceil \alpha \rceil_i$ is the result of replacing in α each proposition $p \in \mathcal{P}_{\overline{EQ_i}}$ by its negation $\neg p$.

For generality, let K be a set of formulas:

Definition 5 *Let $B = (K, U, V)$ be a belief change scenario in $\mathcal{L}_{\mathcal{P}}$ and let $(EQ_i)_{i\in I}$ be the family of all sets of equivalences, as defined in Definition 1.*

Define $\lceil B \rceil$ as $\bigvee_{i\in I}\bigwedge_{(s\in K)}\lceil s \rceil_i$ and $\lceil B \rceil^c$ as $\bigwedge_{(s\in K)}\lceil s \rceil_k$ for selection function c corresponding to EQ_k.

For revision, we define $\lceil (K, \{\alpha\}, \emptyset) \rceil \wedge \alpha$ as the finite representation of $K\dot{+}\alpha$, and analogously $\lceil (K, \{\alpha\}, \emptyset) \rceil^c \wedge \alpha$ as the finite representation of $K\dot{+}_c\alpha$.

Theorem 9 *Let K and α be formulas. Then, we have*

$$\begin{aligned} Cn(K)\dot{+}\alpha &= Cn(\ \lceil (Cn(K), \{\alpha\}, \emptyset) \rceil\ \wedge \alpha) \\ &\equiv \lceil (K, \{\alpha\}, \emptyset) \rceil \wedge \alpha. \end{aligned}$$

Consider example (1): $\{p \wedge q\}\dot{+}(\neg p \vee \neg q)$. So $B = (\{p \wedge q\}, \{(\neg p \vee \neg q)\}, \emptyset)$ is the belief change scenario. We obtain:

$$\lceil B \rceil \wedge (\neg p \vee \neg q) = [(p \wedge \neg q) \vee (\neg p \wedge q)] \wedge (\neg p \vee \neg q),$$

which is equivalent to $(p \equiv \neg q)$. For the other examples in Table 1, if K is the formula corresponding to K' in the first column, then revising by the given α via Theorem 9 is the formula given in the last line (up to permutation of symbols and elimination of definitional equivalents).

Contraction is handled somewhat differently. This is not altogether surprising, given that revision and contraction are not fully interdefinable (Theorem 4). Whereas for revision we replaced each atomic proposition in $\overline{EQ_i}$ by its negation in K, for contraction replacements in K are done over all truth values of atomic propositions in $\overline{EQ_i}$. Formally, given a belief change scenario B, a corresponding set of equivalences EQ_i (according to Definition 1) along with its induced partition $\langle \mathcal{P}_{EQ_i}, \mathcal{P}_{\overline{EQ_i}}\rangle$ of $\mathcal{P}$, and a function $\pi_{k_i} : \mathcal{P}_{\overline{EQ_i}} \to \{\top, \bot\}$, we define for $\alpha \in \mathcal{L}_{\mathcal{P}}$, $\lfloor \alpha \rfloor^{k_i}$ as the result of replacing in α each proposition $p \in \mathcal{P}_{\overline{EQ_i}}$ by $\pi_{k_i}(p)$. Note that each set of equivalences induces a whole set Π_i of such mappings π_{k_i}, viz. $\Pi_i = \{\pi_{k_i} \mid \pi_{k_i} : \mathcal{P}_{\overline{EQ_i}} \to \{\top, \bot\}\}$, amounting to all possible truth assignments to $\mathcal{P}_{\overline{EQ_i}}$.

Definition 6 *Let B and $(EQ_i)_{i\in I}$ be defined as in Definition 5.*

Define $\lfloor B \rfloor$ as $\bigvee_{i\in I, \pi_j\in\Pi_i}\bigwedge_{(s\in K)}\lfloor s \rfloor^j$ and $\lfloor B \rfloor^c$ as $\bigvee_{\pi_j\in\Pi_k}\bigwedge_{(s\in K)}\lfloor s \rfloor^j$ for some selection function c with $c(I) = k$.

We define $\lfloor (K, \emptyset, \{\neg\alpha\}) \rfloor$ as the finite representation of $K\dot{-}\alpha$, and analogously $\lfloor (K, \emptyset, \{\neg\alpha\}) \rfloor^c$ as the finite representation of $K\dot{-}_c\alpha$.

Theorem 10 *Let K and α be formulas. Then, we have*

$$\begin{aligned} Cn(K)\dot{-}\alpha &= Cn(\ \lfloor (Cn(K), \emptyset, \{\neg\alpha\}) \rfloor\ ) \\ &\equiv \lfloor (K, \emptyset, \{\neg\alpha\}) \rfloor\ . \end{aligned}$$

Consider example (2): $\{p \wedge \neg q\} \dot{-} (\neg q)$. We obtain

$$\lfloor (\{p \wedge \neg q\}, \emptyset, \{q\}) \rfloor = (p \wedge \bot) \vee (p \wedge \top) \equiv p\ .$$

For the examples in Table 2, if K is the formula corresponding to K', then in contracting by the given α, the result of the contraction via Theorem 10 is the formula given in the last line (up to permutation of symbols and elimination of definitional equivalents).

Theorems 9 and 10 are interesting in that they show that revision and contraction can be defined with respect to syntactic objects (viz. sentences representing the knowledge base) yet are essentially independent of syntactic form. Hence in a certain sense the approach combines the advantages of base revision (Nebel 1992) and syntax-independent approaches.

Limiting the range of EQ: Intuitively, if an atomic sentence appears in a knowledge base K but not in a sentence for revision α, or vice versa, then that atomic sentence plays no part in the revision. This is indeed the case here, as the next result demonstrates. Let $Vocab(\delta)$ be the atomic sentences in δ. We obtain:

Theorem 11 *Let K be a set of formulas and α a formula. Let $E = Cn(K' \cup EQ \cup \alpha) \cap \mathcal{L}_{\mathcal{P}}$ be a consistent definitional extension of belief change scenario $B = (Cn(K), \{\alpha\}, \emptyset)$.*

Then $\{p \equiv p' \mid p \in (Vocab(K) \setminus Vocab(\alpha)) \cup (Vocab(\alpha) \setminus Vocab(K))\} \subseteq EQ$.

So for belief change, we need consider just the atomic sentences common to K and α, and can ignore (with regards EQ) other atomic sentences. As detailed in the full paper, this result allows one to limit the primed atomic propositions in K' to those occurring in α.

The General Approach

Definition 1 is quite general; in Definitions 2 and 3 we narrow the scope to specific approaches to belief change. We note however, briefly, that other approaches are expressible in this framework. Belief *update* is a distinct form of belief change, suited to a changing world. Update and its dual operator *erasure* are studied in (Katsuno & Mendelzon 1992) where sets of postulates characterising the operators are given.

Definition 7 (Prime Implicate) *A consistent set of literals l is a* prime implicate[1] *of K iff: $l \vdash K$ and for $l' \subset l$ we have $l' \not\vdash K$.*

Definition 8 (Update) *Let K be a knowledge base and α a formula and let $PI(K)$ be the set of prime implicates of K. For each $K_j \in PI(K)$, $1 \leq j \leq m$, let $E^j_1, \ldots, E^j_{n_j}$ be the consistent definitional extensions of $(K_j, \{\alpha\}, \emptyset)$. Then $K \diamond \alpha = \bigcup_{j=1}^{m} \bigcap_{i=1}^{n_j} E^j_i$ is the* update *of K by α.*

We do not define *choice update* here, given space limitations.

[1] Note that this is the dual of *prime implicant*.

Theorem 12 *$K \diamond \alpha$ satisfies the update postulates of (Katsuno & Mendelzon 1992).*

We show in the full paper that the operator $\diamond$ provides a syntactic counterpart for Winslett's update operator (Winslett 1988). We can also take a different notion of *maximal* in Definition 1, and base the definition on set cardinality, rather than set containment. We show that based on this measure we can capture the revision approaches of (Dalal 1988) and (Forbus 1989). Lastly a minor modification to Definition 1 allows one to use the framework to capture the *merging* of knowledge bases.

Conclusion

We have presented a general consistency-based framework for belief change, having the same flavour as the consistency-based paradigms for diagnosis or default reasoning. We focus on a specific approach, in which a knowledge base K and sentence α are expressed, via renaming propositions in K, in separate alphabets. Given this, we assume that as many corresponding atoms in each language are equivalent insofar as is consistently possible. Lastly, we express the resultant knowledge base in a single language. For the revision of K by α, for example, we begin with α and incorporate as much of K as consistently possible. This gives rise to two notions of revision: a choice notion, in which one such "extension" is used for the revised state, and the intersection of all such extensions.

The approach is well-suited for implementation: The notion of a choice extension gives better complexity results than general revision; also, belief revision is an area in which choice reasoning may be useful. Second, we show how the approach can be expressed in terms of a finite knowledge base, and that the scope of a revision can be restricted to those propositions common to the knowledge base and sentence for revision.

The approach allows for a uniform treatment of integrity constraints, in that belief change may take into account both consistency-based and entailment-based static constraints, as well as dynamic constraints. As well, it supports iterated revision. Finally, the framework is applicable to other approaches to belief change.

References

Boutilier, C. 1994. Unifying default reasoning and belief revision in a modal framework. *Artificial Intelligence* 68(1):33–85.

Dalal, M. 1988. Investigations into theory of knowledge base revision. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, 449–479.

Darwiche, A., and Pearl, J. 1997. On the logic of iterated revision. *Artificial Intelligence* 89:1–29.

de Kleer, J., and Konolige, K. 1989. Eliminating the fixed predicates from a circumscription. *Artificial Intelligence* 39(3):391–398.

del Val, A. 1993. Syntactic characterizations of belief change operators. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 540–545.

Eiter, T., and Gottlob, G. 1992. On the complexity of propositional knowledge base revision, updates, and counterfactuals. *Artificial Intelligence* 57:227–270.

Forbus, K. 1989. Introducing actions into qualitative simulation. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1273–1278.

Gärdenfors, P. 1988. *Knowledge in Flux: Modelling the Dynamics of Epistemic States*. Cambridge, MA: The MIT Press.

Katsuno, H., and Mendelzon, A. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence* 52(3):263–294.

Katsuno, H., and Mendelzon, A. 1992. On the difference between updating a knowledge base and revising it. In Gärdenfors, P., ed., *Belief Revision*, 183–203. Cambridge University Press.

Kowalski, R. 1978. Logic for data description. In Gallaire, H., and Minker, J., eds., *Logic and Data Bases*. Plenum Press. 77–103.

Lehmann, D. 1995. Belief revision, revised. In *Proceedings of the International Joint Conference on Artificial Intelligence*, 1534–1540. Montréal: Morgan Kaufmann Publishers.

Liberatore, P., and Schaerf, M. 1997. Reducing belief revision to circumscription (and vice versa). *Artificial Intelligence* 93(1–2):261–296.

Nayak, A. C.; Foo, N. Y.; Pagnucco, M.; and Sattar, A. 1996. Changing conditional belief unconditionally. In *Proc. of the Sixth Conference on Theoretical Aspects of Reasoning About Knowledge*, 119–136.

Nebel, B. 1992. Syntax based approaches to belief revision. In Gärdenfors, P., ed., *Belief Revision*, 52–88. Cambridge University Press.

Poole, D. 1988. A logical framework for default reasoning. *Artificial Intelligence* 36(1):27–48.

Reiter, R. 1984. Towards a logical reconstruction of relational database theory. In Brodie, M.; Mylopoulos, J.; and Schmidt, J., eds., *On Conceptual Modelling*. Springer-Verlag.

Reiter, R. 1987. A theory of diagnosis from first principles. *Artificial Intelligence* 32(1):57–96.

Satoh, K. 1988. Nonmonotonic reasoning by minimal belief revision. In *Proceedings of the International Conference on Fifth Generation Computer Systems*, 455–462.

Spohn, W. 1988. Ordinal conditional functions: A dynamic theory of epistemic states. In Harper, W., and Skyrms, B., eds., *Causation in Decision, Belief Change, and Statistics*, volume II. Kluwer Academic Publishers. 105–134.

Williams, M.-A. 1994. Transmutations of knowledge systems. In Doyle, J.; Torasso, P.; and Sandewall, E., eds., *Proceedings of the Fourth International Conference on the Principles of Knowledge Representation and Reasoning*, 619 – 629.

Winslett, M. 1988. Reasoning about action using a possible models approach. In *Proceedings of the AAAI National Conference on Artificial Intelligence*, 89–93.