Knowledge-Base Revision Using Implications as Hypotheses

Özgür L. Özçep

Institute for Software Systems (STS)
Hamburg University of Technology
Hamburg, Germany
oezguer.oezcep@tu-harburg.de

Abstract. In semantic integration scenarios, the integration of an assertion from some sender into the knowledge base (KB) of a receiver may be hindered by inconsistencies due to ambiguous use of symbols; hence a revision of the KB is needed to preserve its consistency. This paper analyses the new family of implication based revision operators, which exploit the idea of revising hypotheses on the semantic relatedness of the receiver's and sender's symbols. In order to capture the specific inconsistency resolution strategy of these operators, the novel concept of uniform sets, which are based on prime implicates, is elaborated. According to two main results of this paper these operators lend themselves to practical use in systems for semantic integration: First, the operators are finitely representable. Second, the non-sceptical versions of these operators can be axiomatically characterised by postulates, which provide a full specification of the operators' effects.

Keywords: belief revision, semantic integration, postulate, prime implicate

1 Introduction

Belief revision [1] deals with the problem of integrating an assertion stemming from an agent (sender) into a knowledge base (KB) of another agent (receiver). If the receiver trusts the incoming information—and classical belief revision assumes he does—the integration may trigger a revision of the KB because the trigger may be incompatible with the KB; hence some of its formulas have to be eliminated. Belief revision explains the incompatibility with false information in the KB. Therefore, the elimination of formulas in the KB is an adequate means.

But if the diagnosis for the incompatibility is not false information but ambiguous use of symbols, a different strategy seems more appropriate. For example, suppose an agent (the receiver) uses the terminus "article" to denote a publication either in proceedings or in journals while the sender agent uses it to mean publications in journals only. The receiver has different sentences in his KB in which he uses "article" in this sense. So, a trigger sentence stemming from the sender may lead to inconsistencies with the receiver's KB. In order to resolve the inconsistencies, it would not be a good idea to eliminate only one sentence of the KB that contains "article" and that is involved in the conflict; because the next time the receiver integrates a (different) trigger from the sender, the other interpretation of "article" may again lead to inconsistencies.

An appropriate means to deal with conflicts caused by ambiguous use of symbols between different agents is first to state hypotheses on the semantical relatedness of symbols from different agents and second to eliminate some of the hypotheses that are involved in the conflict. This is the general approach of semantic integration based on semantic mappings (or bridging axioms) for heterogeneous knowledge bases [4, 17, 21]. Every KB is assigned a unique name space, and semantic mappings associate symbols of different name spaces. In the case of the example above this means distinguishing between the use of "article" in the receiver's name space and in the sender's name space and initially hypothesising that the uses are equivalent. If the integration of a trigger containing "article" into the receiver's KB leads to inconsistency, a proper strategy for resolving the conflict is eliminating the equivalence hypothesis and possibly replacing it by a weaker hypothesis compatible with the trigger (e.g., by hypothesising that the sender's use is narrower (wider) than the receiver's use).

Based on this strategy for inconsistency resolution, this paper investigates a new class of operators for revising propositional KBs with propositional triggers. The hypotheses used in these operators are implications of the form $p' \rightarrow p$ or $p \rightarrow p'$ where p' stands for the p in the name space of the receiver, and p is the p of the sender. These operators generalise the revision operators of [6] which considers biimplications of the form $p \leftrightarrow p'$ only. Using implications rather than biimplications allows for a more fine-grained analysis of what caused the conflict between the sender's trigger and the receiver's KB.

Though the technical definitions of the revision operators of this paper and of [6] are similar, the theory developed in this paper deviates considerably from that in [6]. On of its main innovative features is a formal specification and analysis of the uniformity property which distinguishes the implication (and biimplication) based operators from classical belief-revision operators. The main idea of the analysis is first to equivalently represent the KB by its most atomic components (prime implicates) and then describe the effect of the implication based operators on the prime implicates by uniform closure conditions.

The implication based revision operators provide a useful abstract implementation model for semantic integration scenarios in which conflicts caused by ambiguous use of symbol between heterogeneous KBs have to be resolved. Though the definitions of the operators are based on infinite sets, they can be described equivalently by finite operators that are appropriate for implementation means (see Th. 2). This is the first main result of this paper. Moreover, anyone implementing the non-sceptical versions of these operators gets a declarative specification of their properties (including uniformity): as a second main result (Th. 4) this paper describes a set of axiomatic postulates which are fulfilled by the operators and which characterise them in the sense that all other operators fulfilling them are representable as implication based choice revision operators.

The paper is structured as follows: The second section provides background on propositional logic and belief revision. The third section discusses the revision operators of J. Delgrande and T. Schaub [6]. The following section introduces the implication based revision operators and shows that these are indeed different from the operators of Delgrande and Schaub. Moreover, the finite representability by a partial polarity flipping operator is proved. The last section before the section on related work and the conclusion gives an axiomatic characterisation of non-sceptical implication based revision operators by postulates.

Proofs of all results in this paper can be found in the technical report [18].

2 Logical Preliminaries

This section introduces notation and concepts from propositional logic and belief revision that are used in the paper. I take for granted the syntax and semantics (interpretation, entailment etc.) of propositional logic, the notion of (sub)clause and the notion of the conjunctive (disjunctive) normal form, CNF (DNF).

Let $\mathcal{P}$ be a set of propositional symbols; $\mathrm{form}(\mathcal{P})$ denotes the set of propositional logical formulas over $\mathcal{P}$, which are denoted by lowercase greek letters $\alpha, \beta \ldots$. Finite sets of formulas are called *knowledge bases* or *belief bases* and are denoted by B as well as primed and indexed variants of B (e.g. B_1, B', $\bar{B}$). $\mathrm{symb}(B)$ is the set of propositional symbols in B. $\mathrm{Int}(\mathcal{P})$ denotes the set of *interpretations* (*assignments*) $\mathcal{I}$ with domain $\mathcal{P}$. $\mathcal{I} \models B$ for a set B is a short notation for $\mathcal{I} \models \bigwedge B$. The set of *consequences* of B over the set of propositional symbols S is $\mathrm{Cn}^S(B) = \{\alpha \in \mathrm{form}(S) \mid B \models \alpha\}$. If the index is left out in some context, then the consequences have to be understood with respect to the maximal set of propositional symbols discussed in the context. If two sets B_1 and B_2 have the same sets of consequences of formulas in $\mathrm{form}(S)$, write $B_1 \equiv_S B_2$. For $\alpha \in \mathrm{form}(\mathcal{P})$ and $S \subseteq \mathcal{P}$, the *clausal closure* of α w.r.t. S is the set $\mathrm{clauseCl}^S(\alpha)$ of clauses that have only symbols from S and that follow from α.

Let Θ_S denote an operator that, given a formula α and a set S of symbols $S \subseteq \mathcal{P}$, computes a formula representing all consequences of α that do not contain symbols in S. (Compare the general framework of forgetting in [14].) This operator will be used as a technical aid for calculating belief-revision results based on hypotheses. For $\mathcal{I} \in \mathrm{Int}(S)$ let $\alpha_{\mathcal{I}}$ be defined as follows: Substitute all occurrences of $p \in S$ in α where $p^{\mathcal{I}} = \mathcal{I}(p) = 1$ by $\top$, else $\bot$ is substituted for p. Now let $\Theta_S : \alpha \mapsto \bigvee_{\mathcal{I} \in \mathrm{Int}(S)} \alpha_{\mathcal{I}}$. For arbitrary $S \subseteq \mathcal{P}$ let $\Theta_S(\alpha) = \Theta_{\mathrm{symb}(\alpha) \cap S}(\alpha)$. For example, let $\alpha = (p \wedge q) \vee (r \wedge s)$ and $S = \{p, r\}$. Then $\Theta_S = ((\bot \wedge q) \vee (\bot \wedge s)) \vee ((\bot \wedge q) \vee (\top \wedge s)) \vee ((\top \wedge q) \vee (\bot \wedge s)) \vee ((\top \wedge q) \vee (\top \wedge s))$. This is equivalent to the formula $s \vee q$. The following facts concerning Θ_S for $S \subseteq \mathcal{P}$ can be proved easily. For all $\alpha \in \mathrm{form}(\mathcal{P})$: $\alpha \models \Theta_S(\alpha)$ and $\mathrm{Cn}^{\mathcal{P} \setminus S}(\alpha) = \mathrm{Cn}^{\mathcal{P} \setminus S}(\Theta_S(\alpha))$. Note, that $\Theta_S(\alpha)$ can be described as the quantified boolean formula $\exists S.\alpha$.

The new operators defined in this paper are based on the concept of *dual remainder sets*, a concept similar to the concept of remainder sets [2] used in the classical paper of Alchourrón, Gärdenfors and Makinson (AGM) [1] for the construction of partial-meet revision functions. Let $B \top \alpha$, *the dual remainder* *sets modulo* α, denote the set of inclusion maximal subsets X of B that are consistent with α, i.e., $X \in B \top \alpha$ iff $X \subseteq B$, $X \cup \{\alpha\}$ is consistent and for all $\bar{X} \subseteq B$ with $X \subset \bar{X}$ the set $\bar{X} \cup \{\alpha\}$ is not consistent. The notion of dual remainders is extended to arbitrary belief bases B_1 as second argument by defining $B \top B_1$ as $B \top \bigwedge B_1$.

An analysis of belief-revision functions involves the investigation of postulates they fulfil. Some postulates for belief-base revision operators $*$ that I will refer to are given below. (In contrast to belief-sets [1] belief bases [10] do not have to be logically closed.)

(BR1) $B * \alpha \not\models \bot$ if $\alpha \not\models \bot$.
(BR2) $\alpha \in B * \alpha$.
(BR3) $B * \alpha \subseteq B \cup \{\alpha\}$.
(BR4) For all $\beta \in B$ either $B * \alpha \models \beta$ or $B * \alpha \models \neg\beta$.
(BR5) If for all $\bar{B} \subseteq B$: $\bar{B} \cup \{\alpha\} \models \bot$ iff $\bar{B} \cup \{\beta\} \models \bot$, then $(B * \alpha) \cap B = (B * \beta) \cap B$.

Postulate (BR1) is the *consistency postulate* [1]; it says that the revision result has to be consistent in case the trigger α is consistent. Postulate (BR2) is the *success postulate* [8]; the revision must be successful in so far as α has to be in the revision result. (BR3) is called the *inclusion postulate* for belief-base revision [12, p. 200]. The revision result of operators fulfilling it are bounded from above. Postulate (BR4) is the *tenacity postulate* [9]; it states that the revision result is complete with respect to all formulas of B. Postulate (BR5) is the *logical uniformity postulate* for belief-base operators [11]. It says that the revision outcomes are determined by the subsets (in)consistent with the trigger.

3 Revision Based on Hypotheses

One example for belief-revision operators that are based on hypotheses are the operators of Delgrande and Schaub [6]. The general idea is to internalize the symbols of the receiver's KB thereby dissociating the name spaces of the receiver, who holds the KB, and the sender, who is the holder of another KB from which the trigger stems. Both name spaces are related by one special form of formula (bridging axiom), namely the biimplication. Holding to a biimplication $p \leftrightarrow p'$ means believing that the propositional symbol p of the receiver (holder of B) has the same meaning as the propositional symbol p of the sender. In order to resolve inconsistencies the internalized KB stays untouched, but some subset of the biimplications are eliminated. If $p \leftrightarrow p'$ is eliminated during the revision process, this can be interpreted as diagnosing the inter-ambiguity of p between the sender and the agent as the culprit for the inconsistency. After the elimination the name space dissociation is abandoned by retaining only those formulas of the old vocabulary. I recapitulate the definitions of the operators and their properties because the revision operator I will introduce is an extension, which uses implications $p' \rightarrow p$ and $p \rightarrow p'$ as hypotheses.

For a given set of propositional symbols $\mathcal{P}$ let $\mathcal{P}'$ denote the set $\{p' \mid p \in \mathcal{P}\}$ of *internal* or *internalized* propositional symbols. Similarly B' denotes the pendant of B where all symbols p are substituted by the corresponding internalized variant p'. I use the following space saving abbreviations (for $p \in \mathcal{P}$): $\overleftrightarrow{p} = p \leftrightarrow p'$, $\overrightarrow{p} = p \rightarrow p'$ and $\overleftarrow{p} = p' \rightarrow p$. A *belief-change scenario* $\langle B_1, B_2, B_3 \rangle$ consists of three sets B_i ($i \in \{1, 2, 3\}$) of formulas over the set of propositional symbols $\mathcal{P}$. B_1 is the initial KB of the receiver, B_2 is a KB that must be contained in the change result and B_3 is a KB that is not allowed to be in the change result. Classical revision of B with α is modelled by the belief-change scenario $\langle B, \{\alpha\}, \emptyset \rangle$; classical contraction of B with α is modelled by the belief-change scenario $\langle B, \emptyset, \{\alpha\} \rangle$. A *belief-change extension* [6, p. 9] (*bc extension* for short) of the belief change scenario $\langle B_1, B_2, B_3 \rangle$ is a set of the form $\mathrm{Cn}^{\mathcal{P}}(B_1' \cup B_2 \cup EQ_i)$, where $EQ_i \subseteq EQ = \{\overleftrightarrow{p} \mid p \in \mathcal{P}\}$ is an inclusion maximal set of biimplications fulfilling the following integrity condition: $\mathrm{Cn}(B_1' \cup B_2 \cup EQ_i) \cap (B_3 \cup \{\bot\}) = \emptyset$. If no such EQ_i exists, then let $\mathrm{form}(\mathcal{P})$ be the only bc extension.

In case of classical belief revision—represented by $\langle B, \{\alpha\}, \emptyset \rangle$—the set of bc extensions E_i have the form $E_i = \mathrm{Cn}^{\mathcal{P}}(B' \cup EQ_i \cup \{\alpha\})$, where $\bot \notin \mathrm{Cn}(B' \cup \{\alpha\} \cup EQ_i)$. Let $(E_i)_{i \in I}$ be the family of all bc extensions in the belief-change scenario $\langle B, \{\alpha\}, \emptyset \rangle$. A selection function c over the index set I selects exactly one index $c(I) \in I$. With these notions, the operators of *choice revision* $\dotplus_c$ based on a selection function c and *sceptical revision* $\dotplus$ are defined as follows.

Definition 1. *[6, p. 11] $B \dotplus_c \alpha = E_k$ (for $c(I) = k$) and $B \dotplus \alpha = \bigcap_{i \in I} E_i$.*

Though the revision results under both operators $\dotplus_c, \dotplus$ are not finite, Delgrande and Schaub can show that these operators are finitely representable. That is more formally, for an operator $\circ \in \{\dotplus_c, \dotplus\}$ one can define an operator $\circ^{fin}$ such that it operates on a finite KB B as left argument, a formula α as right argument and outputs a finite KB $B \circ^{fin} \alpha$ such that $\mathrm{Cn}(B) \circ \alpha = \mathrm{Cn}(B \circ^{fin} \alpha)$. The corresponding finite operators are based on substituting propositional symbols by their negation, thereby flipping the polarity of the symbols. Let $\mathcal{B} = \langle B, \{\alpha\}, \emptyset \rangle$ be a bc scenario and EQ_i a set of biimplications. The formula $\lceil \alpha \rceil_i$ results from α by substituting all occurrences of propositional symbols $p \in \mathcal{P} \setminus \mathrm{symb}(EQ_i)$ with their negation $\neg p$. Let $(E_i)_{i \in I}$ be the family of bc extensions over $\mathcal{B}$ and c a selection function with $c(I) = k$. Then Delgrande and Schaub define the flipping operators by $\lceil \mathcal{B} \rceil = \bigvee_{i \in I} \bigwedge_{\beta \in B} \lceil \beta \rceil_i$ and $\lceil \mathcal{B} \rceil_c = \bigwedge_{\beta \in B} \lceil \beta \rceil_k$ and finite revision operators by:

$$B \dotplus_c^{fin} \alpha = \lceil (B, \{\alpha\}, \emptyset) \rceil_c \wedge \alpha \text{ and } B \dotplus^{fin} \alpha = \lceil (B, \{\alpha\}, \emptyset) \rceil \wedge \alpha$$

The finite representability is stated in Theorem 1.

Theorem 1. *[6, p. 17] $B \dotplus_c \alpha \equiv B \dotplus_c^{fin} \alpha$ and $B \dotplus \alpha \equiv B \dotplus^{fin} \alpha$.*

This theorem evokes a new perspective on what has caused the inconsistency between the KB and the trigger: a flip in the polarity of a propositional symbol. By re-flipping the propositional symbols that caused the inconsistency the result becomes consistent. A remarkable point here is that the flip of a propositional symbol concerns all its occurrences in the formula, it is a kind of uniform flipping. This uniformity can be interpreted as a systematic use of the proposition in just the opposite sense. Referring to the example of the introduction, the receiver would have to substitute all occurrences of “article” by its negation.

4 Using Implications as Hypotheses

By using the set of implications $\mathit{Impl} = \{\overrightarrow{p}, \overleftarrow{p} \mid p \in \mathcal{P}\}$ as set of hypotheses instead of the set of biimplications $EQ = \{\overleftrightarrow{p} \mid p \in \mathcal{P}\}$ new classes of revision operators result. This generalisation from biimplications to implications as hypotheses allows for a more fine-grained diagnosis of the properties of the symbol p that are responsible for the inconsistency. While in the case of Delgrande's and Schaub's operators the diagnosis is a rough “The sender's and the receiver's p have different meanings and so an inconsistency is caused”, the implication based operators account for the “direction” in which the inconsistency was caused. If the hypothesis $p' \rightarrow p$ is eliminated, but $p \rightarrow p'$ is kept in the revision result, then the diagnosis for the inconsistency is the following: the hypothesis that the meaning of the receiver's p is narrower than (or equal to) the sender's meaning of p leads to an inconsistency. But still we can hold on to the hypothesis that the meaning of the sender's p is narrower than the receiver's meaning of p.

The notion of belief extension for biimplication based revision is easily adapted to the case of implications; a set $\mathrm{Cn}^{\mathcal{P}}(B \cup \{\alpha\} \cup X)$ is an *implication based belief extension* iff $X \in \mathit{Impl}\top(B' \cup \{\alpha\})$. (Remember that $\top$ denotes the operator for dual remainder sets defined in Section 2). Let $(\mathit{Impl}_i)_{i \in I}$ be the set of all implication based consistent belief set extensions of $\langle B, \{\alpha\}, \emptyset \rangle$ and c be a selection function for I with $c(I) = k$. The new operators are defined as follows:

Definition 2. *The* implication based choice revision $\dot{+}_c^{\mathit{Impl}}$ *and the* implication based sceptical revision $\dot{+}^{\mathit{Impl}}$ *are defined by :*

$$B \dot{+}_c^{\mathit{Impl}} \alpha = \mathit{Impl}_k \ \textit{(for } c(I) = k\textit{) and } B \dot{+}^{\mathit{Impl}} \alpha = \bigcap_{i \in I} \mathit{Impl}_i$$

As in the case of the Delgrande/Schaub revision operators, one can finitely represent the results by an operation on the KB. The finite representation uses the notion of positive and negative occurrences of propositional symbols. For convenience, I assume that only the connectors $\wedge, \vee$ and $\neg$ are allowed in the formulas; this is no real restriction as this set of connectors is functionally complete. An occurrence of a propositional symbol is *syntactically positive* iff it occurs in the scope of an even number of negation symbols, otherwise it is *syntactically negative*. I also speak of the *(positive, negative) polarity* of a propositional symbol's occurrence. In contrast to the polarity switching of [6], the operator of partial flipping does not change the polarity of all occurrences of a symbol p but only of those of a particular polarity—depending on which implication $\overrightarrow{p}$ or $\overleftarrow{p}$ is missing in the given set of implications.

Definition 3. *Let* $(Impl_i)_{i \in I}$ *be the family of belief extensions for a belief-change scenario* $\mathcal{B} = (B, \{\alpha\}, \emptyset)$ *and let* $Impl_k$ *be an implication based belief extension chosen by the selection function,* $c(I) = k$. *Then define the operator of* partial flipping $\lceil \mathcal{B} \rceil_k^{Impl} = \lceil \mathcal{B} \rceil_c^{Impl}$ *in the following way: If* $p \to p' \notin Impl_k$, *then switch the polarity of the negative occurrences of* p *in* $\bigwedge B$ *(by adding* $\neg$ *in front of these occurrences). If* $p' \to p \notin Impl_k$, *then switch the polarity of the positive occurrences of* p *in* $\bigwedge B$. *Let* $\lceil \mathcal{B} \rceil^{Impl} = \bigvee_{i \in I} \lceil \mathcal{B} \rceil_i^{Impl}$.

With this definition at hand, the following representation theorem follows:

Theorem 2. *The following equivalences hold:*
$B \dot{+}_c^{Impl} \alpha \equiv \lceil (B, \{\alpha\}, \emptyset) \rceil_c^{Impl} \wedge \alpha$ *and* $B \dot{+}^{Impl} \alpha \equiv \lceil (B, \{\alpha\}, \emptyset) \rceil^{Impl} \wedge \alpha$

A simple example shows that $\dot{+}_c^{Impl}$ is different from the operators $\dot{+}_c$, $\dot{+}$.

Example 1. Let be given $\mathcal{P} = \{p, q\}$, $B = \{p \leftrightarrow q\}$, and $\alpha = \neg(p \leftrightarrow q)$. Writing B in CNF (as $(p \vee \neg q) \wedge (\neg p \vee q)$), one can see that it has a positive and a negative occurrence of p, q, respectively. But these different polarities are not dealt with by the biimplication based hypotheses. The two inclusion maximal sets of biimplications are $EQ_1 = \{\overleftrightarrow{p}\}$ and $EQ_2 = \{\overleftrightarrow{q}\}$. Let $I = \{1, 2\}$ and $c_1(I) = 1$, $c_2(I) = 2$. Using $\Theta_{\{p', q'\}}$ or the representation theorem we can calculate the outcomes: $B \dot{+}_{c_1} \alpha = B \dot{+}_{c_2} \alpha = B \dot{+} \alpha = \mathrm{Cn}^{\mathcal{P}}(p \leftrightarrow \neg q)$.

On the other hand, the implication based revision operator recognizes the polarities of the propositional symbols; hence, more possibilities to resolve the conflict result. Here, there are four possibilities given by the following four inclusion maximal sets of implications: $Impl_1 = \mathrm{Cn}^{\mathcal{P}}(\{\overleftrightarrow{q}, \overrightarrow{p}\})$, $Impl_2 = \mathrm{Cn}^{\mathcal{P}}(\{\overleftrightarrow{q}, \overleftarrow{p}\})$, $Impl_3 = \mathrm{Cn}^{\mathcal{P}}(\{\overrightarrow{q}, \overleftrightarrow{p}\})$, and $Impl_4 = \mathrm{Cn}^{\mathcal{P}}(\{\overleftarrow{q}, \overleftrightarrow{p}\})$. These lead to four different choice revisions. Let $I = \{1, 2, 3, 4\}$ and $c(I) = i$. The corresponding revision results are: $B \dot{+}_{c_1}^{Impl} \{\alpha\} = B \dot{+}_{c_4}^{Impl} \{\alpha\} = \mathrm{Cn}^{\mathcal{P}}(\neg p \wedge q)$ and $B \dot{+}_{c_2}^{Impl} \{\alpha\} = B \dot{+}_{c_3}^{Impl} \{\alpha\} = \mathrm{Cn}^{\mathcal{P}}(p \wedge \neg q)$. For illustration, the calculation of the equation $B_1 := B \dot{+}_{c_1}^{Impl} \{\alpha\} = \mathrm{Cn}^{\mathcal{P}}(\neg p \wedge q)$ is given below.

$$
\begin{aligned}
B_1 &= \mathrm{Cn}^{\mathcal{P}}(\{p' \leftrightarrow q', \neg(p \leftrightarrow q), \overleftrightarrow{q}, \overrightarrow{p}\}) \\
&= \mathrm{Cn}^{\mathcal{P}}(\Theta_{\{p', q'\}}((p' \leftrightarrow q') \wedge \neg(p \leftrightarrow q) \wedge \overleftrightarrow{q} \wedge \overrightarrow{p})) \\
&= \mathrm{Cn}^{\mathcal{P}}((\neg(p \leftrightarrow q) \wedge q) \vee (\neg(p \leftrightarrow q) \wedge \neg q \wedge \neg p)) \\
&= \mathrm{Cn}^{P}((\neg(p \leftrightarrow q) \wedge q)) = \mathrm{Cn}^{P}(q \wedge \neg p)
\end{aligned}
$$

In particular, $\dot{+}_{c_1}^{Impl}$ gives results different from those of $\dot{+}_{c_1}, \dot{+}_{c_2}$ and $\dot{+}$.

The example above does not exclude the possibility that the sceptical versions of the biimplication based revision operators and the sceptical versions of the implication based revision operators are the same; it could be the case that the effects of a fine-grained conflict resolving strategy by distinct maximal sets of implications nullify each other. But again, we can show with an example that the use of implications as (enhanced) set of hypotheses has different affects on sceptical revision than the use of biimplications.

Example 2. Let be given $B = (\neg p \wedge q \wedge r \wedge \neg t) \vee (p \wedge \neg q \wedge r \wedge t)$ and $\alpha = (p \wedge \neg q \wedge r \wedge \neg t) \vee (\neg p \wedge q \wedge \neg r \wedge t)$. The maximal sets of biimplications are $EQ_1 = \{\overleftrightarrow{r}, \overleftrightarrow{t}\}$ and $EQ_2 = \{\overleftrightarrow{r}, \overleftrightarrow{p}, \overleftrightarrow{q}\}$. For neither of these sets the model corresponding to $\mathcal{I} := \neg p \wedge q \wedge \neg r \wedge t$ is implied. More concretely, using Theorem 1, one calculates: $B \dot{+} \alpha = (\lceil B \rceil_1 \vee \lceil B \rceil_2) \wedge \alpha = ((p \wedge \neg q \wedge r \wedge \neg t) \vee (\neg p \wedge q \wedge r \wedge t) \vee (\neg p \wedge q \wedge r \wedge t) \vee (p \wedge \neg q \wedge r \wedge \neg t)) \wedge \alpha \equiv (p \wedge \neg q \wedge r \wedge \neg t)$. In contrast to this, there is a maximal set of implications $Impl_1$ that together with $B' \cup \{\alpha\}$ implies $\mathcal{I}$, namely $Impl_1 = \{\overleftrightarrow{t}, \overrightarrow{p}, \overleftarrow{q}, \overrightarrow{r}\}$. So one can calculate:

$$B \dot{+}_1^{Impl} \alpha = ((\neg p \wedge q \wedge \neg r \wedge \neg t) \vee (\neg p \wedge q \wedge \neg r \wedge t)) \wedge \alpha \equiv \neg p \wedge q \wedge \neg r \wedge t$$

Now, $B \dot{+}_1^{Impl} \alpha \models B \dot{+}^{Impl} \alpha$; hence $\mathcal{I} \models B \dot{+}^{Impl} \alpha$ but $\mathcal{I} \not\models B \dot{+} \alpha$.

5 A Representation Theorem

Following the usual approach in classical belief revision [1], I will characterise the non-sceptical implication based revision operators $\dot{+}_c^{Impl}$ by postulates. According to the terminology used in the belief revision literature (cf. [12]), the main theorem of this section (Theorem 4) can be described as a *representation* result: there is a set of postulates such that the class of revision operators $\dot{+}_c^{Impl}$ represents (modulo equivalence) all revision operators fulfilling that set of postulates. Using postulates is a well established methodology in belief revision for declaratively specifying the properties (or the interface) of revision operators that one wants to construct or has constructed. In addition to an implementation-independent specification of revision operators, postulates offer a logical means to compare different revision operators.

The main distinctive feature of Delgrande's and Schaub's operators $\dot{+}_c, \dot{+}$ as well as of $\dot{+}_c^{Impl}, \dot{+}^{Impl}$ is that these operate on a finite set B of formulas as left argument, but do not depend on the specific representation of B. So in contrast to belief-base revision operators they are operators on the knowledge level [16] and thus should be termed *knowledge-base revision operators* [7]. In order to adapt the postulates for belief-base revision one has to replace all references to the set B and its subsets by syntax insensitive concepts.

The key for the adaptation is the use of prime implicates entailed by the KB B. Roughly, prime implicates are the most atomic clauses implied by B. Let be given a set of propositional symbols $\mathcal{P}$ and a subset $S \subseteq \mathcal{P}$ thereof. Let $\alpha \in \text{form}(\mathcal{P})$. Let α be a non-tautological formula. The set $\text{prime}^S(\alpha)$ of prime implicates of α over S is defined in the following way.

$$\text{prime}^S(\alpha) = \{\beta \in \text{clauseCl}^S(\alpha) \mid \emptyset \not\models \beta \text{ and } \beta \text{ has no proper subclause in } \text{clauseCl}^S(\alpha)\}$$

For tautological formulas α let $\text{prime}^S(\alpha) = \{p \vee \neg p\}$, where p is the first propositional symbol occurring in α with respect to a fixed order of $\mathcal{P}$. For example, let $\alpha = (p \vee q) \wedge (\neg q \vee r)$ and $S = \{p, q, r\}$. Then $\text{prime}^{\{p,q,r\}}(\alpha) = \{p \vee q, \neg q \vee r, p \vee r\}$. For knowledge bases let $\text{prime}^S(B) = \text{prime}^S(\bigwedge B)$.

A well known but fundamental fact is that the set of prime implicates of a KB B is equivalent to B itself: $\mathrm{prime}(B) \equiv B$. An additional relevant fact is that if $B_1 \equiv B_2$, then $\mathrm{prime}(B_1) = \mathrm{prime}(B_2)$. These facts justify the perspective on the set of prime implicates as a canonical representation for the knowledge contained in the KB. Moreover, these facts are a useful means for understanding the syntax-insensitive conflict resolution strategy of knowledge-base revision operators.

A second adaptation of the belief-base postulates concerns the uniformity of the operators $\dot{+}_c, \dot{+}$ as well of $\dot{+}_c^{Impl}, \dot{+}^{Impl}$. The conflicts between B and the trigger α are handled on the level of symbols and not on the level of formulas. Therefore, in order to mirror this effect on the prime implicates one has to impose a uniformity condition. If, e.g., the hypothesis $p' \rightarrow p$ is eliminated in the conflict resolution process, then formulas of the knowledge base B, in which p occurs positively, are not preserved in the revision result. In general, if a set of implication based hypotheses Im is given, then $B' \cup Im$ preserves a subset of prime implicates of B which fulfils some closure condition concerning the polarities of symbols. These sets of prime implicates can be characterised as uniform sets according to the following definition.

Definition 4. *Let $B \subseteq \mathrm{form}(\mathcal{P})$ be a KB. A set $X \subseteq \mathrm{prime}(B)$ is called* uniform w.r.t. to B and implications, *$X \in U^{Impl}(B)$ for short, iff the following closure condition holds: If $pr \in \mathrm{prime}(B)$ is such that (a) $\mathrm{symb}(pr) \subseteq \mathrm{symb}(X)$ and (b) for all symbols p in pr there is a $pr_p \in X$ that contains p in the same polarity, then pr is contained in X, i.e., $pr \in X$.*

Example 3. Let $B = \{p \vee q, p \vee r \vee s, r \vee t, s \vee u\}$. Then $\mathrm{prime}(B) = B$. Now, among all subsets $X \subseteq \mathrm{prime}(B)$ only the set $X := \{p \vee q, r \vee t, s \vee u\}$ is not uniform as it would have to contain $p \vee r \vee s$, too. Formally, $U^{Impl}(B) = \mathrm{Pow}(\mathrm{prime}(B)) \setminus \{\{p \vee q, r \vee t, s \vee u\}\}$. ($\mathrm{Pow}(X)$ denotes the power set of X, i.e. the set of all subsets of X.)

A proper justification for Definition 4—in the sense that it really captures the intended concept—is Theorem 3 below. It shows that for all B, Im one can find a uniform set X that is equivalent to $B' \cup Im$. The set X exactly describes the collection of logical atoms (prime implicates) of the receiver's KB B that are preserved after dissociating the name spaces of the sender and receiver (step from B to B') and adding hypotheses on the semantical relatedness in Im.

Theorem 3. *Let $\mathcal{P}$ and $\mathcal{P}'$ be disjoint sets of propositional symbols. Let B be a KB and σ be a injective substitution for some subset $S = \{p_1, \ldots, p_n\} \subseteq \mathcal{P}$ such that $\sigma(S) = \{p_1', \ldots, p_n'\} \subseteq \mathcal{P}'$ and let Im be a set of implication based hypotheses containing at most primed symbols of $\sigma(S)$. Then there is a uniform set $X \in U^{Impl}(B)$ such that: $B' \cup Im \equiv_{\mathcal{P}} X$.*

Now, we give postulates for revision operators $*$ that characterise the implication based choice revision operators. They are variants of the postulates mentioned in the section on logical preliminaries.

(R1) $B * \alpha \not\models \bot$ if $B \not\models \bot$ and $\alpha \not\models \bot$.

(R2) $B * \alpha \models \alpha$.
(R3) There is a set $H \subseteq U^{Impl}(B)$ s.t. $B * \alpha \equiv \bigwedge \bigcup H \wedge \alpha$ or $B * \alpha \equiv \bigwedge \bigcup H$.
(R4) For all $X \in U^{Impl}(B)$ either $B * \alpha \models X$ or $B * \alpha \models \neg \bigwedge X$.
(R5) For all $Y \subseteq U^{Impl}(B)$: If $\bigcup Y \cup \{\alpha\} \models \perp$ iff $\bigcup Y \cup \{\beta\} \models \perp$, then $\{X \in U^{Impl}(B) \mid B * \alpha \models X\} = \{X \in U^{Impl}(B) \mid B * \beta \models X\}$.

Postulate (R1) can be termed the postulate of weak consistency; it says that the revision result has to be consistent (satisfiable) in case both the trigger α and the KB B are consistent. The consistency postulate for AGM belief revision and belief base revision (BR1) is stronger as it demands the consistency also in the case where only α is consistent. Postulate (R2) is a weak success postulate; the revision must be successful in so far as the result has to imply α. It is weaker than the postulate (BR2) for belief bases. (R3) is an adapted version of the inclusion postulate for belief base revision (BR3), which can be rewritten as: There is a $\bar{B} \subseteq B$ such that $B * \alpha = \bar{B} \cup \{\alpha\}$ or $B * \alpha = \bar{B}$. In (R3) B is replaced by the set of uniform sets w.r.t. B, and set identity is shifted to equivalence. Postulate (R4) can be called uniform tenacity. It is a very strong postulate, which states that all uniform sets w.r.t. to B either follow from the result or are falsified. This postulate captures the maximality of the operator $\dot{+}_c^{Impl}$. Postulate (R5) is an adaptation of the logical uniformity postulate for belief-base operators (BR5). It says that the revision outcomes w.r.t. to the revision operator $*$ are determined by the uniform sets implied by the revision result.

Postulates (R1)–(R5) are sufficient to represent the class of implication based choice revision operators modulo equivalence.

Theorem 4. *A revision operator $*$ fulfils the postulates (R1)–(R5) iff it can be equivalently described as $\dot{+}_c^{Impl}$ for some selection function c.*

6 Related Work

The work described in this paper follows in general the belief-revision tradition as initiated by the pioneering work of AGM [1], but has main differences due to a different explanation of the inconsistencies. Moreover, classical belief-revision functions à la AGM operate on a logically closed (and hence infinite) set called *belief set* and a formula which triggers the revision of the belief set into a new belief set. In belief-base revision [10] the revised KB is allowed to be an arbitrary not necessarily closed (finite) set of sentences called *belief base*. The negative property of belief-base revision of being syntax sensitive is remedied in the case of *knowledge-base revision operators* which are exemplified by the revision operators of this paper as well as those of [6] and [5].

The revision operators of this paper are based on the elimination of hypotheses that have the role of semantic mappings [17]. The idea of using belief revision techniques to revise semantic mappings has already been worked in the literature [15], [21]. But these approaches consider the set of semantic mappings as the object of revision, while the approach of this paper considers the semantic mappings as revision aids that are deleted after the revision.

The notion of a prime implicate is used in the approaches of [20], [22], [3]. In contrast to the approach of this paper, these do not use prime implicates in the formulation of the postulates; they (only) define new belief-revision operators based on prime implicates and show that they fulfil some classical postulates.

The implication based revision operators exhibit a symbol-oriented rather than a sentence-oriented strategy for inconsistency resolution. A different symbol-oriented approach is described by Lang and Marquis [13]. Their revision operators do not use hypotheses but the well-known concept of forgetting [14].

7 Conclusion and Outlook

I have presented a new type of revision operator, which resulted as a generalisation of Delgrande's and Schaub's operators [6] by considering implications rather than biimplications as hypotheses. Similar to a result of [6], it can be shown that the operators are finitely representable and hence suitable for implementation.

But we have seen (cf. beginning of Section 4) that the generalisation from biimplications to implications adds the value of having a more fine-grained diagnosis of what exactly leads to the ambiguity. Moreover, I described postulates that integrate the uniformity property in order to characterise the implication based operators. Delgrande and Schaub [6] show which (classical) postulates their operators fulfil but do not give a representation theorem. In this paper, I could at least show that the implication based choice revision can be characterised by a set of postulates (Theorem 4). A slightly different notion of uniform set leads to a representation theorem for biimplication based choice revision.

I motivated the perspective to consider the sets of biimplications and implications as hypotheses on the semantical relatedness of symbols belonging to different name spaces. This perspective leads naturally to the question what other initial sets of hypotheses on the semantical relatedness could be used as a basis for new revision operators. In fact, one could consider bridging axioms like $p' \leftrightarrow q$, which relate symbols hypothesised to be synonyms. Using a set H of such creative hypotheses may induce operators that are quite different from classical revision operators as the former may not be conservative: $B' \cup H$ may imply formulas $\beta \in \text{form}(\mathcal{P})$ that do not already follow from B. Such creative behaviour does not occur for $H = Im_i$ or $H = EQ_i$.

There already exist approaches in the area of ontology alignment where more expressive semantic mappings are handled (e.g., [4]). We note that the framework of this paper is extendable to more expressive KR formalisms like first order logic [19] by using a more syntactical notion of prime implicate.

References

1. Alchourrón, C.E., Gärdenfors, P., Makinson, D.: On the logic of theory change: partial meet contraction and revision functions. Journal of Symbolic Logic 50, 510–530 (1985)

2. Alchourrón, C.E., Makinson, D.: Hierarchies of regulations and their logic. In: Hilpinen, R. (ed.) New Studies in Deontic Logic, pp. 125–148. D. Reidel Publishing (1981)
3. Bienvenu, M., Herzig, A., Qi, G.: Prime implicate-based belief revision operators. In: Ghallab, M., Spyropoulos, C.D., Fakotakis, N., Avouris, N.M. (eds.) ECAI. vol. 178, pp. 741–742. IOS Press (2008)
4. Bouquet, P., Giunchiglia, F., van Harmelen, F., Serafini, L., Stuckenschmidt, H.: Contextualizing ontologies. Web Semantics: Science, Services and Agents on the World Wide Web 1(4), 325–343 (2004)
5. Dalal, M.: Investigations into a theory of knowledge base revision: preliminary report. In: Proceedings of the 7th National Conference on Artificial Intelligence (AAAI-88). pp. 475–479. AAAI Press, St. Paul, Minnesota (August 1988)
6. Delgrande, J.P., Schaub, T.: A consistency-based approach for belief change. Artificial Intelligence 151(1–2), 1–41 (2003)
7. Eschenbach, C., Özçep, Ö.L.: Ontology revision based on reinterpretation. Logic Journal of the IGPL 18(4), 579–616 (2010), first published online August 12, 2009
8. Gärdenfors, P.: Rules for rational changes of belief. In: Pauli, T. (ed.) Philosophical Essays Dedicated to Lennart Aquist on his Fiftieth Birthday, pp. 88–101. Philosophical Society and Department of Philosophy, Uppsala University (1982)
9. Gärdenfors, P.: Knowledge in Flux: Modeling the Dynamics of Epistemic States. The MIT Press, Bradford Books, Cambridge, MA (1988)
10. Hansson, S.O.: Belief Base Dynamics. Ph.D. thesis, Uppsala University (1991)
11. Hansson, S.O.: Reversing the Levi identity. J. of Phil. Logic 22, 637–669 (1993)
12. Hansson, S.O.: A Textbook of Belief Dynamics. Kluwer Academic Publishers (1999)
13. Lang, J., Marquis, P.: Reasoning under inconsistency: A forgetting-based approach. Artificial Intelligence 174(12-13), 799–823 (2010)
14. Lin, F., Reiter, R.: Forget it! In: In Proceedings of the AAAI Fall Symposium on Relevance. pp. 154–159 (1994)
15. Meilicke, C., Stuckenschmidt, H.: Reasoning support for mapping revision. Journal of Logic and Computation (2009)
16. Newell, A.: The knowledge level. Artificial Intelligence 18, 87–127 (1982)
17. Noy, N.F.: Semantic integration: a survey of ontology-based approaches. SIGMOD Record 33(4), 65–70 (2004)
18. Özçep, O.L.: Knowledge-base revision using implications as hypotheses (extended version). Technical report, Institute for Softwaresystems (STS), Hamburg University of Technology (2012), available online at `http://www.sts.tu-harburg.de/tech-reports/papers.html` or `http://dl.dropbox.com/u/65078815/oezcep12KnowledgeTRExtended.pdf`
19. Özçep, O.L.: Minimality postulates for semantic integration. Accepted for publication in the proceedings of BNC@ECAI2012 (2012)
20. Pagnucco, M.: Knowledge compilation for belief change. In: Sattar, A., Kang, B.h. (eds.) AI 2006: Advances in Artificial Intelligence, LNCS, vol. 4304, pp. 90–99. Springer Berlin / Heidelberg (2006)
21. Qi, G., Ji, Q., Haase, P.: A conflict-based operator for mapping revision. In: Grau, B.C., Horrocks, J., Motik, B., Sattler, U. (eds.) Proceedings of the 22nd International Workshop on Description Logics (DL-09). CEUR Workshop Proceedings, vol. 477 (2009)
22. Zhuang, Z., Pagnucco, M., Meyer, T.: Implementing iterated belief change via prime implicates. In: Orgun, M., Thornton, J. (eds.) AI 2007: Advances in Artificial Intelligence, LNCS, vol. 4830, pp. 507–518. Springer Berlin / Heidelberg (2007)