

Belief Revision Revisited

Ewa Madalińska-Bugaj¹ and Witold Łukaszewicz²

¹ Institute of Informatics, Warsaw University, Warsaw, Poland

² Dept. of Computer Science, Linköping University, Sweden and
College of Economics and Computer Science TWP, Olsztyn, Poland

Abstract. In this paper, we propose a new belief revision operator, together with a method of its calculation. Our formalization differs from most of the traditional approaches in two respects. Firstly, we formally distinguish between defeasible observations and indefeasible knowledge about the considered world. In particular, our operator is differently specified depending on whether an input formula is an observation or a piece of knowledge. Secondly, we assume that a new observation, but not a new piece of knowledge, describes exactly what a reasoning agent knows at the moment about the aspect of the world the observation concerns.

1 Introduction

Belief revision [1] is the task of modifying a reasoner's knowledge base when new information becomes available. More formally, given a knowledge base KB , representing the reasoner's belief set, and a piece of new information α , the task is to specify the new reasoner's knowledge base $KB * \alpha$. There are three important assumptions underlying belief revision. Firstly, it is supposed that the reasoner's knowledge base is incomplete and possibly incorrect. Secondly, the reasoner's environment is assumed to be static.¹ Thirdly, whenever a new piece of information is inconsistent with the current knowledge base, new information is considered more reliable than the knowledge base.²

The classical specification of belief revision has been proposed in [1] in the form of eight *rationality postulates*, known in the AI literature as AGM postulates. Two of them are of special interest in this paper.

(R1) If $KB \not\models \neg\alpha$, then $KB + \alpha \subseteq KB * \alpha$, where $KB + \alpha$ is the deductive closure of $KB \cup \{\alpha\}$.

(R2) If $KB_1 \equiv KB_2$, then $KB_1 * \alpha \equiv KB_2 * \alpha$.

The next example shows that the postulate (R1) is sometimes very problematic from the intuitive point of view.

Example 1. Watching TV yesterday, I learned that on the next Sunday there would be rain in Paris. So my knowledge base KB is $\{r\}$. Watching TV today, I have learned that on the next Sunday there will be rain or snow in Paris, i.e. $\alpha = r \vee s$. According

¹ There is another important form of belief change, called *belief update* [2]. In contrast to belief revision, it deals with dynamic settings, where a piece of new information is the result of a performed action.

² A comprehensive literature on the subject of belief revision can be found in [3].

to (R1), the resulting knowledge base, $KB * \alpha$, should contain r . However, intuition dictates that $KB * \alpha = Cn(r \vee s)$.³ ■

Note that when we say that the resulting knowledge base in Example 1 should be $Cn(r \vee s)$, we make an implicit assumption that a new observation is exactly what an agent knows at the moment about the aspect of the world the observation is concerned with. Thus, if a new observation is weaker than KB , KB should be weakened.⁴

Consider now the postulate (R2). At the first glance, it seems to be indisputable. However, as the following example illustrates, the situation is more subtle.

Example 2. Let $KB_1 = \{p, p \Rightarrow s\}$ and $\alpha = \neg p$, where p and s stand for “Tweety is a penguin” and “Tweety is a bird”, respectively. Since the truth of s directly depends on the truth of p , intuition dictates that the resulting knowledge base is $KB * \alpha = Cn(\neg p)$.⁵

Consider now the knowledge base $KB_2 = \{p, s\}$ and $\alpha = \neg p$, where p and s stand for “Mr Smith is rich” and “Mr Jones is rich”, respectively. In this case, $KB * \alpha$ should be $Cn(\neg p \wedge s)$. On the other hand, KB_1 and KB_2 are logically equivalent. ■

Although KB_1 and KB_2 from Example 2 are logically equivalent, they differ significantly as regards the type of information they contain. Whereas facts like “Tweety is a bird” or “Mr Jones is rich” represent an agent’s observations about the considered world, the sentence “If Tweety is a penguin, Tweety is a bird” represents rather the agent’s knowledge about the world that can be used to draw conclusions from observations.⁶

Example 2 shows that we should distinguish between observations and a general knowledge about the world under consideration. And we should treat this knowledge as more reliable than an ordinary observation.

In this paper, we propose a new formalization of belief revision. It differs from the traditional approaches in two respects. Firstly, it is always assumed that new information describes exactly what a reasoning agent knows at the moment about the aspect of the world the observation concerns. Secondly, we formally distinguish between observations and knowledge about the considered world. More specifically, a knowledge base is not a set of formulae, but a pair of sets of formulae, $\langle OB, A \rangle$, where OB and A represent observations of an agent and its knowledge (i.e. domain axioms) about the world, respectively. Whereas observations have status of beliefs, i.e. they can be invalidated by a piece of new information, formulae representing the agent’s knowledge about the world are assumed to be always true.

Domain axioms correspond closely to integrity constraints considered in the theory of data (knowledge) bases. However, there is a subtle difference between these notions.

³ Here Cn stands for the consequence operator of classical propositional logic.

⁴ In [4], we are presented with a new form of belief revision, called *conservative belief change*, where a similar assumption is made. The relationship between this approach and our proposal presented in the rest of this paper will be discussed in section 6.

⁵ One can also argue that the resulting knowledge base should be represented in an equivalent form, namely $Cn(\neg p \wedge (p \Rightarrow s))$. This will make it possible to retrieve s , if the next piece of new information is p again.

⁶ The term “observation” here means either an observation made directly by an agent or communicated to the agent by other sources.

Integrity constraints are usually assumed to be fixed and external with respect to a data base. Domain axioms, on the other hand, are considered as a part of a knowledge base and a new domain axiom can be learned by a reasoning agent. As we shall see later, $KB * \alpha$ should be differently specified depending on whether α is an observation or a new domain axiom.

The paper is structured as follows. In section 2, we provide preliminary definitions. In section 3, we formally describe our belief revision operator under the assumption that an input formula is a new observation. We also illustrate our proposal by considering a number of examples. In section 4, we specify our belief revision operator under the assumption that an input formula is a new domain axiom. In section 5, we shortly discuss our proposal in the context of AGM postulates. Section 6 is devoted to related work. Finally, section 7 contains concluding remarks and future work.

2 Preliminaries and Terminology

We deal with a propositional language with a finite set of propositional symbols, called atoms. We assume that each language under consideration contains two special atoms \top and \perp , standing for truth and falsity, respectively. Formulae are built in the usual way using standard connectives \wedge , \vee , \Rightarrow , \neg and \Leftrightarrow .

A formula of the form of p or $\neg p$, where p is an atom, is called a *literal*. Interpretations are identified with maximal consistent sets of literals. For any formula α , we write $|\alpha|$ to denote the set of all *models* of α . We use the symbol Cn to denote the consequence relation of classical propositional logic.

Let α be a formula. By $ATM(\alpha)$ we denote a set of all non-redundant atoms occurring in α . An atom p occurring in α is said to be redundant iff $\alpha[p \leftarrow \top] \equiv \alpha[p \leftarrow \perp] \equiv \alpha$ ⁷.

Let p be an atom and suppose that α is a formula. We write $\exists p.\alpha$ to denote the formula $\alpha[p \leftarrow \top] \vee \alpha[p \leftarrow \perp]$. If $P = \{p_1, \dots, p_n\}$ is a set of atoms and α is a formula, then $\exists P.\alpha$ stands for $\exists p_1 \dots \exists p_n.\alpha$.

A formula of the form $\exists P.\alpha$, where $P = \{p_1, \dots, p_n\}$, is called an *eliminant* of $\{p_1, \dots, p_n\}$ in α . Intuitively, such an eliminant can be viewed as a formula representing the same knowledge as α about all atoms from $ATM(\alpha) - P$ and providing no information about the atoms in P . Formally, this property is stated by the following theorem [5].

Theorem 1. Let α and β be formulae such that

$$ATM(\beta) \subseteq (ATM(\alpha) - P), \text{ where } P = \{p_1, \dots, p_n\}.$$

Then $\alpha \models \beta$ iff $(\exists P.\alpha) \models \beta$. ■

A *clause* is a formula of the form $l_1 \vee \dots \vee l_n$, $n \geq 1$, where l_i , $1 \leq i \leq n$, is a literal.

We say that a clause c' *absorbs* a clause c if c' is a subclause⁸ of c . For instance, the clause a absorbs the clause $a \vee l$. Let α be a formula in conjunctive normal form (CNF).

⁷ $\alpha[p \leftarrow \top]$ (resp. $\alpha[p \leftarrow \perp]$) is the formula obtained from α by replacing all occurrences of p by \top (resp. \perp).

⁸ A clause c' is a subclause of c iff c' entails c , but not vice versa.

We write $ABS(\alpha)$ to denote the formula obtained from α by deleting all absorbed clauses. Clearly, α and $ABS(\alpha)$ are equivalent.

Two clauses are said to have an *opposition* if one of them contains a literal l and the other the literal $\neg l$.

Suppose that two clauses, c_1 and c_2 , have exactly one opposition. Then the *resolvent* of c_1 and c_2 , written $res(c_1, c_2)$, is the clause obtained from the disjunction $c_1 \vee c_2$ by deleting the opposed literals as well as any repeated literals. For example, $res(\neg a \vee l, a \vee d)$ is $l \vee d$.

Definition 1. Let α be a formula. We say that a clause c is a prime implicate of α iff

- (i) $\alpha \Rightarrow c$ is a tautology;
- (ii) there is no clause c' which is a subclause of c and $\alpha \Rightarrow c'$ is a tautology. ■

Algorithm 2. Let α be a formula. The *prime implicates form of α* , written $PIF(\alpha)$, is the formula obtained from α by the following construction.

1. Let β be the conjunctive normal form of α .
2. Repeat as long as possible:
if β contains a pair c and c' of clauses whose resolvent exists and no clause of β is a subclause of $res(c, c')$, then $\beta := \beta \wedge res(c, c')$.
3. Take $ABS(\beta)$. This is $PIF(\alpha)$. ■

The following result holds ([6]).

Theorem 2. Let α be a formula.

- (i) $PIF(\alpha)$ is a conjunction of all prime implicates of α .
- (ii) $PIF(\alpha)$ and α are equivalent.
- (iii) All atoms occurring in $PIF(\alpha)$ are non-redundant. ■

Let α and β be formulae. A *tail* of α and β , written $TL(\alpha, \beta)$, is the conjunction of those prime implicates of $\alpha \wedge \beta$ that are neither prime implicates of α nor prime implicates of β . Intuitively, $TL(\alpha, \beta)$ can be viewed as those additional conclusions which can be derived by combining α and β . The formula $TL(\alpha, \beta)$ can be constructed using the following algorithm.

Algorithm 3.

1. $\gamma := PIF(\alpha) \wedge PIF(\beta)$; $TL(\alpha, \beta) := \top$.
2. Repeat as long as possible:
if γ contains a pair c and c' of clauses whose resolvent exists and no clause of γ is a subclause of $res(c, c')$, then
 $\gamma := \gamma \wedge res(c, c')$;
 $TL(\alpha, \beta) := TL(\alpha, \beta) \wedge res(c, c')$.
3. $TL(\alpha, \beta) := ABS(TL(\alpha, \beta))$. ■

3 Defining Belief Revision Operator

Definition 4. A knowledge base is a pair $KB = \langle OB, A \rangle$, where OB is a finite set of formulae, called observations, and A is a finite set of formulae, called domain axioms. In the sequel, we will never distinguish between finite sets of formulae and their conjunctions. In particular, both OB and A will be often considered as single formulae. Any knowledge base KB uniquely determines a belief set. This set, denoted by $Cn(KB)$, is the set $Cn(OB \wedge A)$. KB is said to be consistent iff $Cn(KB)$ is consistent. ■

Now we define our revision operator $*$. There are two cases to consider.

1. New information is an ordinary observation.
2. New information is a new piece of knowledge.

In this section we consider the former of the above cases. The latter will be discussed in section 4.

We start with some intuitions underlying our approach. Suppose that $KB = \langle OB, A \rangle$ is a consistent knowledge base and α is a new observation. Recall that we make an implicit assumption that a new observation is exactly what is known at the moment about the aspect of the world it concerns. In particular, if a new observation is weaker than KB , KB should be suitably weakened. A natural way to achieve this goal is to delete all information concerning atoms occurring in α . This can be technically done using eliminants⁹. Denote the weakened observation formula by OB' . The formula $OB' \wedge \alpha$ is a natural candidate for the observation formula in the revised knowledge base $KB * \alpha$. It is easily seen that $OB' \wedge \alpha$ and $OB' \wedge A$ are both consistent. The problem, however, is that $OB' \wedge \alpha \wedge A$ may be not. If this is the case, OB' should be further weakened. Observe that the source of inconsistency can be new conclusions which can be derived from $A \wedge \alpha$. Note that these new conclusions are represented by $TL(A, \alpha)$. Therefore, we must delete information concerning atoms occurring in those conjuncts of $TL(A, \alpha)$ which are inconsistent with OB' . The resulting formula, strengthened by α , is the observation formula in the revised knowledge base $KB * \alpha$.

The above intuitions are formalized below.

Definition 5. Let $KB = \langle OB, A \rangle$ be a knowledge base and α be a new observation. The new knowledge base, $KB * \alpha$, is $\langle OB_1, A \rangle$, where OB_1 obtains from OB by the following construction.

Let $TL(A, \alpha) = c_1 \wedge \dots \wedge c_n$.

- (1) $P := ATM(\alpha)$;
- (2) $OB' := \exists P.OB$;
- (3) $R := \{\}$;
- (4) for $i := 1$ to n do
if $OB' \wedge c_i \equiv \perp$ then $R := R \cup ATM(c_i)$;
- (5) $OB_1 := \alpha \wedge \exists R.OB'$. ■

⁹ Note that we cannot weaken KB by weakening A , because domain axioms represent knowledge about the world and hence cannot be invalidated.

3.1 Examples

The following examples illustrate an application of Definition 5 to compute $KB * \alpha$.

Example 3. Let $KB = \langle OB, A \rangle$ be a knowledge base, where $OB = p$ and $A = (p \Rightarrow s)$. Suppose further that a new observation is $\alpha = \neg p$. The new knowledge base is $KB_1 = \langle OB_1, A \rangle$, where OB_1 is computed in the following steps:¹⁰

- $ATM(\alpha) = \{p\}$.
- $OB' = \exists p.p \equiv \top$.
- $PIF(A) = \neg p \vee s$; $PIF(\alpha) = \neg p$.
- $TL(A, \alpha) = \top$.
- Since, after performing step (4), $R = \{\}$,
 $OB_1 = \alpha \wedge \top \equiv \neg p$.

Thus $Cn(KB * \alpha) = Cn(\neg p \wedge (p \Rightarrow s))$. Note that $Cn(KB * \alpha)$ does not contain s . ■

Example 4. Let $KB = \langle OB, A \rangle$ be a knowledge base, where $OB = p \wedge s$ and $A = \{\}$. A new observation is $\alpha = \neg p$. The resulting knowledge base is $KB_1 = \langle OB_1, A \rangle$ where OB_1 is computed as follows.

- $ATM(\alpha) = \{p\}$.
- $OB' = \exists p.p \wedge s \equiv s$.
- $PIF(A) = \top$; $PIF(\alpha) = \neg p$.
- $TL(A, \alpha) = \top$.
- $OB_1 = \alpha \wedge OB' = \neg p \wedge s$.

Therefore $Cn(KB * \alpha) = Cn(\neg p \wedge s)$. Observe that according to our intuitions, s is a member of $Cn(KB * \alpha)$. ■

Example 5. Let $KB = \langle OB, A \rangle$ be a knowledge base, where $OB = p$ and $A = (p \Rightarrow s)$ and let a new information α be $\neg s$. We compute OB_1 .

- $ATM(\alpha) = \{s\}$.
- $OB' = \exists s.p \equiv p$.
- $PIF(\alpha) = \neg s$; $PIF(A) = \neg p \vee s$.
- $TL(A, \alpha) = \neg p$.
- After performing step (4), we get $R = \{p\}$.
- $OB_1 = \neg s \wedge \exists p.p \equiv \neg s$.

Thus, $Cn(KB * \alpha) = Cn(\neg s \wedge (p \Rightarrow s))$. Note that according to our intuitions p does not belong to $Cn(KB * \alpha)$. ■

Example 6. Let $KB = \langle OB, A \rangle$ be a knowledge base, where $OB = p \wedge s$ and $A = (p \Rightarrow q) \wedge (q \Rightarrow r)$ and let a new observation α be $\neg q$. The computation of OB_1 is the following.

- $ATM(\alpha) = \{q\}$.
- $OB' = \exists q.p \wedge s \equiv p \wedge s$.

¹⁰ We use the symbol ' \equiv ' as the meta symbol denoting that two formulae are equivalent.

- $PIF(\alpha) = \neg q$; $PIF(A) = (\neg p \vee q) \wedge (\neg q \vee r) \wedge (\neg p \vee r)$.
- $TL(A, \alpha) = \neg p$.
- After performing step (4), we get $R = \{p\}$.
- $OB_1 = ((\exists p.OB') \wedge \alpha) \equiv s \wedge \neg q$

$Cn(KB * \alpha) = Cn(s \wedge \neg q \wedge (p \Rightarrow q) \wedge (q \Rightarrow r))$. Observe that $\neg p$ is a member of $Cn(KB * \alpha)$. ■

4 Absorbing New Knowledge

In this section we define revision operator under the assumption that an input formula α is a new piece of knowledge. To fix some intuitions, consider the following example.

Example 7. Let $KB = \langle \{b\}, \{\} \rangle$, where b stands for “Tweety is a black penguin”. Suppose that I learned that penguins are always black or grey. This allows me to conclude that Tweety is black or grey, so a new piece of knowledge is $b \vee g$, where g stands for “Tweety is a grey penguin”. Clearly, the resulting knowledge base should be $\langle \{b\}, \{b \vee g\} \rangle$ because learning that all penguins are black or grey I should not assume that my earlier observation that Tweety is black should be weakened. ■

The above example illustrates that our definition of belief operator given in section 3 should be modified in the case when new information is a piece of knowledge. Even, as in this paper, if we assume that a new observation always gives us exact information about the aspect of the world it concerns, this assumption makes little sense in the context of a new domain axiom. The role of domain axioms is to put some constraints on the world under consideration, so the only observations that should be invalidated by domain axioms are those that are inconsistent with them.

Suppose that $KB = \langle OB, A \rangle$ is a consistent knowledge base and α is a new domain axiom. Denote the resulting knowledge base KB_1 by $\langle OB_1, A_1 \rangle$. Obviously, A_1 should be $A \wedge \alpha$. Now we would like to put $OB_1 := OB$. The problem, however, is that $OB \wedge A_1$ can be inconsistent. Since the original knowledge base KB is consistent, the source of inconsistency can be only α , i.e. conjuncts from $PIF(\alpha)$, and these new conclusions which can be derived from A and α , i.e. conjuncts from $TL(A, \alpha)$. Combining them together we receive $ABS(TL(A, \alpha) \wedge PIF(\alpha))$. Therefore, we must delete information concerning atoms occurring in those conjuncts of $ABS(TL(A, \alpha) \wedge PIF(\alpha))$ which are inconsistent with OB .

A formalization of above idea is given below.

Definition 6. Let $KB = \langle OB, A \rangle$ be a knowledge base and α be a new piece of knowledge. $KB * \alpha = \langle OB_1, A_1 \rangle$, where $A_1 = A \wedge \alpha$ and OB_1 is obtained from OB by the following construction.

Let $ABS(TL(A, \alpha) \wedge PIF(\alpha)) = c_1 \wedge \dots \wedge c_n$.

- (1) $P := \{\}$;
- (2) for $i := 1$ to n do
if $OB \wedge c_i \equiv \perp$ then $P := P \cup ATM(c_i)$;
- (3) $OB_1 := \exists P.OB$. ■

Notice that the above algorithm is a slight modification of steps (3)-(5) of algorithm presented in Definition 5. Note also that the resulting knowledge base $KB * \alpha$ is inconsistent if and only if the formula $A \wedge \alpha$ is inconsistent.

4.1 Examples

We now present a number of examples illustrating the construction from Definition 6.

Example 8. Let $KB = \langle OB, A \rangle$, where $OB = p \wedge q \wedge s$ and $A = (p \Rightarrow r)$. Suppose that $\alpha = (q \Rightarrow \neg s)$. The new knowledge base $KB * \alpha$ is $\langle OB_1, A_1 \rangle$, where $A_1 = (p \Rightarrow r) \wedge (q \Rightarrow \neg s)$ and OB_1 is computed as follows.

- $ABS(TL(A, \alpha) \wedge PIF(\alpha)) = (\neg q \vee \neg s)$.
- $OB \wedge (\neg q \vee \neg s) \equiv \perp$, so $OB_1 := \exists q, s. OB \equiv p$.

Thus, $Cn(KB * \alpha) = Cn(p \wedge (p \Rightarrow r) \wedge (q \Rightarrow \neg s))$. ■

Example 9. Let $KB = \langle OB, A \rangle$, where OB is $(p \vee q) \wedge (s \vee r)$ and A is $p \Rightarrow r$. Assume that $\alpha = (q \Rightarrow \neg s)$. The new knowledge base $KB * \alpha$ is $\langle OB_1, A_1 \rangle$, where $A_1 = (p \Rightarrow r) \wedge (q \Rightarrow \neg s)$ and OB_1 is computed as follows.

- $ABS(TL(A, \alpha) \wedge PIF(\alpha)) = (\neg q \vee \neg s)$.
- $OB \wedge (\neg q \vee \neg s) \not\equiv \perp$, so $OB_1 := OB$.

Thus, $Cn(KB_1) = Cn((p \vee q) \wedge (s \vee r) \wedge (p \Rightarrow r) \wedge (q \Rightarrow \neg s)) \equiv Cn((p \vee q) \wedge r \wedge (\neg q \vee \neg s))$. ■

Example 10. Let $KB = \langle OB, A \rangle$, where OB is $p \wedge s \wedge (q \vee r)$ and A is $p \Rightarrow r$. Let $\alpha = (r \Rightarrow \neg s)$. The new knowledge base $KB * \alpha$ is $\langle OB_1, A_1 \rangle$, where $A_1 = (p \Rightarrow r) \wedge (r \Rightarrow \neg s)$ and OB_1 is computed as follows.

- $ABS(TL(A, \alpha) \wedge PIF(\alpha)) = (\neg r \vee \neg s) \wedge (\neg p \vee \neg s)$.
- $OB \wedge (\neg r \vee \neg s) \not\equiv \perp$; $OB \wedge (\neg p \vee \neg s) \equiv \perp$. So, $OB_1 = \exists p, s. OB \equiv (q \vee r)$.

Thus, $Cn(KB_1) = Cn((q \vee r) \wedge (p \Rightarrow r) \wedge (r \Rightarrow \neg s))$. ■

5 Postulates

In this section, we specify postulates for our revision operator. We start with some terminology and notation.

Definition 7. Let $KB_1 = \langle OB_1, A_1 \rangle$ and $KB_2 = \langle OB_2, A_2 \rangle$ be knowledge bases. We say that KB_1 and KB_2 are equivalent, denoted by $KB_1 \equiv KB_2$, iff $OB_1 \equiv OB_2$ and $A_1 \equiv A_2$. ■

If X and α are formulae, then $X + \alpha$ stands for $Cn(X \wedge \alpha)$. If $KB = \langle OB, A \rangle$ is a knowledge base and α is a formula, then $KB + \alpha$ stands for $Cn(OB \wedge A \wedge \alpha)$. We write $KB \models \alpha$ iff $(OB \wedge A) \models \alpha$. KB_{OB} and KB_A denote the sets OB and A , respectively.

The following postulates, corresponding loosely to AGM postulates, hold for our revision operator.

- (R1) $KB * \alpha$ is a belief set, i.e. $Cn(KB * \alpha) = KB * \alpha$.
- (R2) $\alpha \in KB * \alpha$.
- (R3) $KB * \alpha \subseteq KB + \alpha$.
- (R4) If $KB \not\models \neg\alpha$ and $KB + \alpha \subseteq Cn((\exists ATM(\alpha).KB_{OB}) \wedge \alpha)$, then $KB + \alpha \subseteq KB * \alpha$.
- (R5) $KB * \alpha$ is inconsistent iff $\{\alpha\} \cup KB_A$ is inconsistent.
- (R6) If $KB_1 \equiv KB_2$, then $KB_1 * \alpha \equiv KB_2 * \alpha$.
- (R7) If $ATM(\alpha) \subseteq ATM(\alpha \wedge \beta)$ and $KB \not\models \neg\alpha \vee \neg\beta$ then $KB * (\alpha \wedge \beta) \subseteq (KB * \alpha) + \beta$.
- (R8) If $KB \not\models \neg\alpha \vee \neg\beta$ and $(KB * \alpha) + \beta \subseteq (\exists ATM(\alpha \wedge \beta).KB_{OB}) + (\alpha \wedge \beta)$, then $(KB * \alpha) + \beta \subseteq KB * (\alpha \wedge \beta)$.

The postulates (R1)-(R3) and (R6) are exactly the AGM postulates, whereas the remaining postulates are weaker forms of the AGM postulates.

6 Related Work

An important property of our formalization of belief revision is the assumption that a new observation is exactly what an agent knows at the moment about the aspect of the world the observation is concerned with. As we remarked earlier, this assumption is also made in [4], where an interesting formalization of belief revision, called *conservative belief revision* (CBR, for short), is presented. However there are two important differences between CBR and our formalization.

- (i) The semantics for CBR is based on Grove's system of spheres ([7]), originally developed for AGM-revision. What typifies Grove's semantics is that the revision operator it defines depends on an ordering over all interpretations (of the considered language) signifying their level of plausibility. In consequence, CBR is not a single belief revision operator, but rather a class of such operators. The problem, of course, is that it is not clear which of them should be chosen in practical applications. Our formalization, on the other hand, provides a unique belief revision operator.
- (ii) In contrast to our approach, CBR does not distinguish between defeasible observations and knowledge about the considered world. As we argued earlier, such distinction is important, because observations are subject to invalidation, whereas knowledge is not.

In [8], a belief update operator, called MPMA, has been defined. As the belief revision operator specified here, MPMA is heavily influenced by the notion of an eliminant. However, there is a crucial difference between these formalisms due to the general difference between belief revision and belief update. As we stated earlier, belief revision is based on the assumptions that a world being modelled is static and beliefs describing a reasoner's environment may be incorrect. In belief update, on the other hand, we assume that a piece of new information represents an effect of a performed action and the current set of the reasoner's beliefs is correct (see [2]). This distinction manifests clearly in the presence of domain axioms (called integrity constraints in MPMA). The next example illustrates this.

Example 11. Let $KB = \langle OB, A \rangle$, where $OB = w$ and $A = (w \Rightarrow a)$ (here a and w stand for “a turkey is alive” and “a turkey is walking”). Thus, $Cn(KB)$ contains $w \wedge a$. Assume that a piece of new information α is $\neg w$. If KB is considered from the belief revision perspective, the resulting knowledge base should be $\neg w$, because there is no reason to believe a when w has turned out to be false. On the other hand, if KB is considered from the update perspective, the resulting knowledge base should be $\neg w \wedge a$, because there is no reason to conclude that an action that have made the turkey non-walking made it dead. ■

7 Conclusions and Future Work

We have presented a new belief revision operator. Our approach assumes that a new observation provides exact information about the aspect of the considered world it concerns. Also, we formally distinguish between defeasible observations and indefeasible knowledge about the considered world.

There are three topics that we left for further research.

- (1) Our belief revision operator has been specified syntactically. It would be interesting to provide its semantical characterization.
- (2) A dual notion to belief revision is belief contraction. This is a task of specifying a new knowledge base under the assumption that some beliefs are retracted but no new beliefs are added. It seems that the notion of an eliminant provides a very natural basis to solve this task.
- (3) AGM postulates do not address the issue of iterated belief revision [9]. On the other hand, in practical applications we are interested in a sequence of belief revisions rather, than in a single one. It would be interesting to investigate our belief operator in this context.

References

1. Alchourrón, C.E., Gärdenfors, P., Makinson, D.: On the logic theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic* **50** (1985) 510–530
2. Katsuno, H., Mendelzon, A.O.: On the difference between updating a knowledge base and revising it. In: *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*. (1991) 387–394
3. Delgrande, J.P., Schaub, T.: A consistency-based approach for belief change. *Artificial Intelligence Journal* **151** (2003) 1–41
4. Delgrande, J.P., Nayak, A.C., Pagnucco, M.: Gricean belief change. *Studia Logica* (2004) To appear. Electronic version can be found at <http://www.cs.sfu.ca/~jim/publications.html>.
5. Brown, F.M.: *Boolean Reasoning*. Kluwer Academic Publishers (1990)
6. Quine, W.: A way to simplify truth functions. *American Mathematical Monthly* **62** (1955) 627–631
7. Grove, A.: Two modellings for theory change. *Journal of Philosophical Logic* **17** (1988) 157–170
8. Doherty, P., Łukaszewicz, W., Madalinska-Bugaj, E.: The PMA and relativizing minimal change for action update. *Fundamenta Informaticae* **44** (2000) 95–131
9. Darwiche, A., Pearl, J.: On the logic of iterated revision. *Artificial Intelligence Journal* **89** (1997) 1–29