Proceedings of the Federated Conference on
Computer Science and Information Systems pp. 81–87
ISBN 978-83-60810-51-4

Revising Structured Knowledge Bases

Michał Korpusik and Witold Łukaszewicz
Faculty of Mathematics and
Computer Science
University of Warmia and Mazury
Olsztyn, Poland
Email: m.korpusik@matman.uwm.edu.pl
witold.lukaszewicz@gmail.com

Ewa Madalińska-Bugaj
Institute of Informatics
Warsaw University
Warsaw, Poland
Email: ewama@mimuw.edu.pl

***Abstract*—In this paper we present a new approach to belief revision. In contrast to traditional formalizations of this problem, where all pieces of information included in a knowledge base have identical status, we explicitly distinguish between observations, i.e., facts which an agent observes or is being told, and facts representing a general, sometimes defeasible, knowledge about the considered world.**[1]

This distinction allows us to deal with scenarios that cannot be properly modelled using existing belief revision operators.

Keywords: Knowledge representation, belief revision

I. Introduction

BELIEF revision [1], [13] is the task of modifying a reasoner's knowledge base when new information becomes available. More specifically, given a knowledge base KB, representing the reasoner's belief set, and a piece of new information α, the task is to determine the new reasoner's knowledge base $KB * \alpha$.[2]

There are three important assumptions underlying belief revision. First, it is assumed that the reasoner's knowledge base KB may be incorrect. Second, the reasoner's environment is supposed to be static.[3] Finally, whenever the new piece of information α is inconsistent with the current knowledge base, new information is considered to be more reliable than the knowledge base.

In [1], we are provided with a number of postulates, usually referred to as the AGM postulates, which any reasonable belief revision operator should satisfy. Katsuno and Medelzon reformulated these postulates for the case when KB is a finite knowledge base. Probably the most uncontroversial of their postulates is the following.

If $KB \equiv KB'$ and $\alpha \equiv \alpha'$, then $KB * \alpha \equiv KB' * \alpha'$.

The above postulate is indisputable. We must not accept different results if equivalent knowledge bases are revised by equivalent formulae. However, consider the following example.

Example 1: Suppose that I observed that Tweety is a bird and I know that birds typically fly. Thus my knowledge base is $KB = \{p, p \Rightarrow q\}$, where p and q stand for "Tweety is a bird" and "Tweety flies", respectively. Suppose further that my new observation α is $\neg p$. Clearly removing p from KB, I have also to remove q.[4] Hence, $KB * \alpha = \{\neg p, p \Rightarrow q\}$.

Assume now that I have been told that Mary is married and John is rich. Hence my knowledge base $KB' = \{p, q\}$, where p and q stand for "Mary is married" and "John is rich", respectively. Let $\alpha = \neg p$. Intuition dictates that $KB' * \alpha$ is $\{\neg p, q\}$. On the other hand KB and KB' are logically equivalent!
■

Example 1 shows that we have to distinguish between formulae representing observations and formulae representing general knowledge that can be used to draw conclusions from observations.[5]

As regards Example 1, in KB $p \Rightarrow q$ represents knowledge, whereas p is an observation. In KB' both p and q are observations.[6]

In this paper we propose a new form of belief revision. It differs from traditional approaches in that observations are separated from formulae representing general knowledge about the world. In addition, the latter formulae are divided into two classes: defeasible statements like "a bird typically flies" and domain axioms like "a penguin is a bird".

The paper is organized as follows. In section II, we provide preliminary definitions. In section III, we present a general belief change framework, defined in [3], which forms a basis for our approach. Our belief revision operator is introduced in section IV. In section V, we shortly discuss our proposal in the context of the AGM postulates. In section VI, we illustrate the introduced operator by considering a number of examples. In

[1] In our earlier paper [10], we postulated to separate observations and general knowledge. However, an approach presented here significantly differs from that proposed in [10]. (See discussion in section IX.)

[2] Introducing a belief revision paradigm, Alchourron et al. assumed that KB is a *belief set*, i.e. a deductively closed set of formulae. However it seems to be more practical to regard KB as a finite set of formulae, representing a belief set ([8]).

[3] That is, α is not the result of an action. The task of revising a knowledge base by a formula being an effect of an action is known in the AI literature as *belief update* ([7], [11], [4]).

[4] This is because our belief that Tweety flies was based on the assumption that she was a bird.

[5] In the philosophical literature the term "knowledge" stands for statements that are true. Here we also use this term to denote defeasible statements like "birds typically fly". For more detailed discussion of this point see Remark 2.

[6] The term "observation" mean here either an observation made directly by an agent or communicated to the agent by other sources.

section VII, we introduce a prioritized version of our operator. In section IX, we provide a comparison with related work. Finally, section X contains concluding remarks.

II. PRELIMINARIES

We deal with propositional language based on a finite set of propositional letters, referred to as *atoms*. We distinguish two special atoms $\top$ and $\bot$, denoting the truth and the falsity respectively. Formulae are constructed in the usual way using standard connectives $\neg$, $\wedge$, $\vee$, $\Rightarrow$ and $\Leftrightarrow$. The set of all formulae over a set of atoms $\mathcal{P}$, called the language over $\mathcal{P}$, will be denoted by $\mathcal{L}_\mathcal{P}$. A *literal* is an atom or its negation.

We write Th to denote the provability operator of propositional logic.We say that a set of formulae $X \subseteq \mathcal{L}_\mathcal{P}$ is deductively closed iff $Th(X) = X$.

A *formula corresponding to* a finite set of formulae X is the conjunction of all members of X. If X is the empty set, then a formula corresponding to X is $\top$.

Let X be any finite set of formulae. We write $X^\wedge$ to denote the conjunction of all members of X.

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

A formula of the form $\exists P.\alpha$, where $P = \{p_1, \ldots, p_n\}$, is called an *eliminant of* $\{p_1, \ldots, p_n\}$ *in* α. Intuitively, such an eliminant can be viewed as a formula representing the same knowledge as α about all atoms occurring in α except those in P and providing no information about the atoms in P. The reader interested in a detailed theory of eliminants should consult Brown [2].[7]

The following theorem holds.

Theorem 1: Let X be a finite set of formulae over $\mathcal{L}_{\mathcal{P}\cup\mathcal{Q}}$, where P and Q are disjoint sets of atoms. Then $Th(X) \cap \mathcal{L}_\mathcal{P} = Th(\exists Q.X^\wedge)$.

■

III. CONSISTENCY-BASED BELIEF CHANGE

Our approach is heavily influenced by a consistency-based framework for expressing belief change, introduced in [3]. This framework is based on the following notion.

Definition 1: A *belief change scenario* is a triple $B = \langle K, R, C\rangle$, where K, R and C are sets of formulae over a fixed propositional language $\mathcal{L}_\mathcal{P}$.

Informally, K is a knowledge base which is to be modified in such a way that the resulting knowledge base includes all elements from R and does not include any element from C. The modified knowledge base corresponding to B will be denoted by $K \dotplus R \dot{-} C$.

■

Remark 1: If C is empty, then a belief change scenario describes belief revision. If R is empty and C contains a single formula, then a belief change scenario describes contraction.[8]

[7] In the AI literature, eliminants are sometimes referred to as *forget operators*, see [9].

[8] Contraction is a dual notion of revision, where beliefs are retracted but no new beliefs are added. This notion will be of no interest here.

■

A central notion connected with a belief change scenario is that of an extension.

Definition 2: Let $B = \langle K, R, C\rangle$ be a belief change scenario over $\mathcal{L}_\mathcal{P}$. Define a new set $\mathcal{P}'$ of atoms, isomorphic with $\mathcal{P}$, given by $\mathcal{P}' = \{p' : p \in \mathcal{P}\}$. Let K' be a knowledge base obtained from K by replacing any $p \in \mathcal{P}$ by $p' \in \mathcal{P}'$. Let EQ be a maximal (wrt set inclusion) set of equivalences $\{p \Leftrightarrow p' | p \in \mathcal{P}\}$ such that $Th(K' \cup EQ \cup R) \cap (C \cup \bot) = \emptyset$. The set $Th(K' \cup EQ \cup R) \cap \mathcal{L}_\mathcal{P}$ is called a *belief change extension* of B.

If there is no such set EQ, then B is inconsistent and $\mathcal{L}_\mathcal{P}$ is a unique (inconsistent) belief change extension of B.

■

Note that EQ can be equivalently defined as a maximal set of equivalences $\{p \Leftrightarrow p' | p \in \mathcal{P}\}$ such that

$$K' \cup EQ \cup R \nvdash \alpha, \text{ for every } \alpha \in C \cup \bot. \quad (1)$$

Intuitively, a belief change extension of B is a modification of K which includes all formulae from R and which includes no formula from C.

For a given belief change scenario, there may be more than one extension. In such a case the modified knowledge base will be identified with the intersection of all extensions.

Definition 3: Let $(E_i)_{i\in I}$ be the class of all belief change extensions of $B = \langle K, R, C\rangle$. Then

$$K \dotplus R \dot{-} C = \bigcap_{i\in I} E_i.$$

■

Example 2: Consider a belief change scenario $B = \langle K, R, C\rangle$, where

$$K = \{p \wedge q \wedge r\}; \quad R = \{\neg p \vee \neg q\}; \quad C = \{r\}.$$

$K' = \{p' \wedge q' \wedge r'\}$. We first look for maximal sets $EQ \subseteq \{p \Leftrightarrow p', q \Leftrightarrow q', r \Leftrightarrow r'\}$ such that $\{p' \wedge q' \wedge r'\} \cup \{\neg p \vee \neg q\} \cup EQ$ is consistent and $\{p' \wedge q' \wedge r'\} \cup \{\neg p \vee \neg q\} \cup EQ \nvdash r$. These are $EQ_1 = \{p \Leftrightarrow p'\}$ and $EQ_2 = \{q \Leftrightarrow q'\}$. Therefore, there are two extensions

$$E_1 = Th(\{p' \wedge q' \wedge r'\} \cup \{\neg p \vee \neg q\} \cup \{p \Leftrightarrow p'\}) \cap L_{\{p,q,r\}}$$

$$E_2 = Th(\{p' \wedge q' \wedge r'\} \cup \{\neg p \vee \neg q\} \cup \{q \Leftrightarrow q'\}) \cap L_{\{p,q,r\}}$$

In conclusion, $K \dotplus R \dot{-} C = E_1 \cap E_2 = Th(\{p \Leftrightarrow \neg q\})$.

■

According to Definitions 2 and 3, a modified knowledge base corresponding to a given belief change scenario is deductively closed and hence, infinite. However, if K, R and C are all finite, a modified knowledge base can be given a finite representation. The details follow.

If an atom appears in a knowledge base K but not in $R \cup C$, then this atom has no influence on extensions of a belief change scenario. Similarly, if an atom appears in $R \cup C$, but not in K. Therefore, while constructing EQ sets, we can restrict our attention to atoms that appears in $K \cap (R \cup C)$.

If X is any set of formulae, then $ATM(X)$ denotes the set of all atoms occurring in X.

The following result can be found in [3].

Theorem 2: Let $B = \langle K, R, C\rangle$ be a belief change scenario over $\mathcal{L}_\mathcal{P}$. Let $Q = ATM(K) \cap (ATM(R) \cup ATM(C))$. Define a new set Q' of atoms, isomorphic with Q, given by $Q' = \{q' : q \in Q\}$. Let K' be a knowledge base obtained from K by replacing any $q \in Q$ by $q' \in Q'$. Let EQ be a maximal (wrt set inclusion) set of equivalences $\{q \Leftrightarrow q' : q \in Q\}$ such that $Th(K' \cup EQ \cup R) \cap (C \cup \bot) = \emptyset$. The set $Th(K' \cup EQ \cup R) \cap \mathcal{L}_\mathcal{P}$ is a *belief change extension* of B.

■

A belief change scenario $B = \langle K, R, C\rangle$ is said to be *finite* if the sets K, R and C are all finite.

The following algorithm can be used for computing extensions of finite belief change scenarios.[9]

Algorithm 1: Let $B = \langle K, R, C\rangle$ be a finite belief change scenario over $\mathcal{L}_\mathcal{P}$. Let Q, Q' and K' be as in Theorem 2.

(1) [Computing EQ sets].
Compute all maximal (wrt set inclusion) sets, $EQ_1, \ldots EQ_n$, of equivalences $\{q \Leftrightarrow q' : q \in Q\}$ such that $Th(K' \cup EQ_i \cup R) \cap (C \cup \bot) = \emptyset$. If there is no such set, then return $\mathcal{L}_\mathcal{P}$ as a sole (inconsistent) extension of B and stop. Otherwise:

(2) [Eliminating primed variables]
Let $\gamma_1, \ldots, \gamma_n$ be formulae corresponding to sets $K' \cup EQ_1 \cup R, \ldots, K' \cup EQ_n \cup R$. Put $FE_i = \exists Q'.\gamma_i$.
Return $Th(FE_1), \ldots, Th(FE_n)$ as extensions of B.

■

The correctness of the above algorithm follows from Theorems 1 and 2.

Since, for any formulae $\alpha_1, \ldots, \alpha_n$, $Th(\alpha_1) \cap \cdots \cap Th(\alpha_n) = Th(\alpha_1 \vee \cdots \vee \alpha_n)$, the following result holds.

Theorem 3: Let $B = \langle K, R, C\rangle$ be a finite belief change scenario and let $FE_1, \ldots, FE_n$ be formulae specified in Algorithm 1. The modified knowledge base corresponding to B is given by

$$K \dotplus R \dot{-} C = Th(FE_1 \vee \cdots \vee FE_n).$$

■

Formulae $FE_1, \ldots, FE_n$ can be viewed as finite representations of extensions of a belief change scenario $B = \langle K, R, C\rangle$, whereas $FE_1 \vee \cdots \vee FE_n$ can be viewed as a finite representation of a modified knowledge base corresponding to the scenario. This disjunction will be denoted by $K \oplus R \ominus C$.

In the sequel we will make use of the following lemma.

Lemma 1: Let $B = \langle K, \{\alpha\}, \{\neg\beta\}\rangle$ be a belief change scenario. If $K \cup \{\alpha\} \cup \{\beta\}$ is consistent, then $K \oplus \{\alpha\} \ominus \{\neg\beta\}$ is equivalent to $K \cup \{\alpha\}$.

Proof Follows immediately from the fact that there is one EQ set associated with B, containing all equivalences of the form $p \Leftrightarrow p'$, where p is any atom from $ATM(K) \cap (ATM(\{\alpha\}) \cup ATM(\{\beta\})$.

■

[9]This algorithm is simpler than a similar algorithm provided in [3].

Example 3 (Continuation of Example 2): Reconsider a belief change scenario from Example 2

$$K = \{p \wedge q \wedge r\}; \quad R = \{\neg p \vee \neg q\}; \quad C = \{r\}.$$

As we saw, there are two EQ sets, namely $EQ_1 = \{p \Leftrightarrow p'\}$ and $EQ_2 = \{q \Leftrightarrow q'\}$. Hence, $\gamma_1 = p' \wedge q' \wedge r' \wedge (p \Leftrightarrow p') \wedge (\neg p \vee \neg q)$ and $\gamma_2 = p' \wedge q' \wedge r' \wedge (q \Leftrightarrow q') \wedge (\neg p \vee \neg q)$. Eliminating p', q' and r', we get $FE_1 = p \wedge (\neg p \vee \neg q) \equiv p \wedge \neg q$ and $FE_2 = q \wedge (\neg p \vee \neg q) \equiv q \wedge \neg p$. Thus

$$K \dotplus R \dot{-} C = Th(FE_1 \vee FE_2) \equiv Th(p \Leftrightarrow \neg q)$$

$$K \oplus R \ominus C = (FE_1 \vee FE_2) \equiv p \Leftrightarrow \neg q.$$

■

IV. Our belief revision operator

In this section we define our belief revision operator.

As we mentioned earlier, we will distinguish between formulae representing observations and those representing general knowledge about the world under consideration. In addition, the latter formulae will be divided into two classes: those representing defeasible knowledge, like "A bird typically flies", and *domain axioms*, representing general facts that are always true.[10]

Remark 2: In the philosophical literature the term "knowledge" is usually reserved for facts that are guaranteed to be true. In particular, the statement "I know α" implies that α is true. On the other hand, there is a weaker notion, namely that of "belief". The statement "I believe α" does not imply that α must be true. It implies only that I am prepared to act as if α were true. (see [14]).

Accordingly, the term "defeasible knowledge", we use here, may seem a little strange. However, in our opinion, a statement like "A bird typically flies" represents a kind of knowledge rather than a belief. This is because it is true that a typical bird flies. Of course, if we use this statement to conclude that a particular bird flies, the conclusion may turn out to be false. This is the reason why the word "knowledge" has been preceded by the word "defeasible".[11]

■

Definition 4: A *knowledge base* is a triple $KB = \langle OB, DS, DA\rangle$, where OB, DS and DA are finite sets of formulae. These sets are referred to as *observations*, *defeasible statements* and *domain axioms*, respectively.

■

Let X be any finite set of formulae. Recall that $X^\wedge$ stands for the conjunction of all members of F.

Definition 5: Let $KB = \langle OB, DS, DA\rangle$ be a knowledge base and suppose that α is a revision formula representing a new observation. A revision of KB by α, written $KB * \alpha$,

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

[11]The term "defeasible knowledge" has been first used in [12].

is a new knowledge base given by $\langle OB_1, DS, DA\rangle$, where $OB_1 = OB \oplus \{\alpha\} \ominus \{\neg DA^\wedge\}$. Here $OB \oplus \{\alpha\} \ominus \{\neg DA^\wedge\}$ is a finite representation of the modified knowledge base corresponding to belief change scenario $\langle OB, \{\alpha\}, \{\neg DA^\wedge\}\rangle$. (See section III.)

■

Observe that defeasible statements and domain axioms are identical in both knowledge bases. The definition of OB_1 requires an explanation. A new set of observations has to satisfy two conditions. First, it should contain α. Second, it must be consistent with DA. The framework presented in the previous section gives us a convenient tool to achieve both of these goals.

Each knowledge base determines a belief set.

Definition 6: Let $KB = \langle OB, DS, DA\rangle$ be a knowledge base. A belief set corresponding to KB, written B_{KB}, is given by $DS \dotplus (OB \cup DA)$.[12]

■

The following result follows immediately from Definitions 5 and 6.

Theorem 4: Let $KB = \langle OB, DS, DA\rangle$ be a knowledge base.

$$B_{KB*\alpha} = DS \dotplus [(OB \oplus \{\alpha\} \ominus \{\neg DA^\wedge\}) \cup DA].$$

■

It is easy to see that the proposed revision operator consists of two belief changes. First, we alter the observation set OB with α while maintaining domain axioms. Next we revise the defeasible pieces of knowledge using the previous result.

Example 4 (Continuation of Example 1): Reconsider first $KB = \langle\{p\}, \{p \Rightarrow q\}, \{\top\}\rangle$ and $\alpha = \neg p$. We first compute $OB_1 = \{p\} \oplus \{\neg p\} \ominus \{\bot\} = \{p\} \oplus \{\neg p\}$. To this end, we look for a maximal set $EQ \subseteq \{p \Leftrightarrow p'\}$ such that $\{p'\} \cup \{\neg p\} \cup EQ \nvdash \bot$. There is one such set, namely the empty set. Thus, the formula corresponding to $\{p'\} \cup \{\neg p\} \cup \{\top\}$, is $\gamma = p' \wedge \neg p$. Eliminating p' from γ, we conclude that $OB_1 = \neg p$. Thus, the revised knowledge base, $KB * \alpha$, is $\langle\{\neg p\}, \{p \Rightarrow q\}, \{\top\}\rangle$. $B_{KB*\alpha} = DS \dotplus (OB_1 \cup \{\top\}) = \{p \Rightarrow q\} \dotplus \{\neg p\} = Th(\{p \Rightarrow q, \neg p\}) = Th(\{\neg p\})$.

Now consider $KB = \langle\{p, q\}, \{\top\}, \{\top\}\rangle$ and $\alpha = \neg p$. It is easily computed that the revised knowledge base $KB * \alpha$ is $\langle\{\neg p, q\}, \{\top\}, \{\top\}$ and $B_{KB*\alpha} = Th(\{\neg p, q\})$.

These results agree with our intuition.

■

V. Properties of our revision operator

Alchourron et al.[1] proposed eight postulates, known in the AI literature as AGM postulates, which, as they claimed, should be satisfied by every revision operator. In this section we analyze our revision operator in the context of these postulates.

We shall need the following notion. Let $KB = \langle OB, DS, DA\rangle$ and $KB' = \langle OB', DS', DA'\rangle$ be two knowledge bases. KB and KB' are said to be equivalent, written $KB \equiv KB'$, iff $OB \equiv OB'$, $DS \equiv DS'$ and $DA \equiv DA'$.

If $KB = \langle OB, DS, DA\rangle$, then $KB + \alpha$ stands for $\langle OB \cup \{\alpha\}, DS, DA\rangle$. We write OB_{KB}, DS_{KB} and DA_{KB} to denote a set of observations, a set of defeasible statements and a set of domain axioms of KB, respectively.

The AGM postulates are the following.[13]

(1) $B_{KB*\alpha}$ is deductively closed.
(2) $\alpha \in B_{KB*\alpha}$.
(3) $B_{KB*\alpha} \subseteq B_{KB+\alpha}$.
(4) If $\neg\alpha \notin B_{KB}$, then $B_{KB+\alpha} \subseteq B_{KB*\alpha}$.
(5) $B_{KB*\alpha}$ is inconsistent iff α is inconsistent.
(6) If $KB \equiv KB'$ and $\alpha \equiv \alpha'$, then $B_{KB*\alpha} \equiv B_{KB'*\alpha'}$.
(7) $B_{KB*(\alpha\wedge\beta)} \subseteq B_{(KB*\alpha)} \cup \beta$.
(8) If $\neg\beta \notin B_{KB*\alpha}$, then $B_{(KB*\alpha)} \cup \beta \subseteq B_{KB*(\alpha\wedge\beta)}$.

Theorem 5: Our belief revision operator satisfies the postulates (1), (2), (3), (4), (6) and (7). The The postulate (5) and (8) are not satisfied.

Proof: The satisfiability of (1), (2) (6) and (7) is obvious.

As regards (3), consider two cases. Assume first that $OB_{KB} \cup \{\alpha\} \cup DA_{KB}$ is inconsistent. Then $B_{KB+\alpha}$ is also inconsistent and we are done. Otherwise, i.e., if $OB_{KB} \cup \{\alpha\} \cup DA_{KB}$ is consistent, (3) holds in view of Lemma 1. The postulate (4) follows immediately from Lemma 1.

To see that (5) does not hold, take $KB = \langle\{\bot\}, \{\top\}, \{\top\}\rangle$ and $\alpha = p$. Clearly, $KB * \alpha$ is inconsistent, whereas α is consistent. It is easily verified that a weaker version of (5) holds, namely $B_{KB*\alpha}$ is inconsistent iff α is inconsistent or $DA \cup \{\alpha\}$ is inconsistent or OB is inconsistent or DS is inconsistent.

The following is a counterexample to (8).[14]

$OB = (p \wedge q \wedge r \wedge s) \vee (\neg p \neg q \wedge \neg r \neg s)$.
$\alpha = (\neg p \wedge \neg q \wedge r \wedge s) \vee (p \wedge \neg q \wedge \neg r \wedge \neg s) \vee (\neg p \wedge \neg q \wedge r \wedge \neg s)$.
$\beta = (\neg p \wedge \neg q \wedge r \wedge s) \vee (p \wedge \neg q \wedge \neg r \wedge \neg s)$.
$DS = \top$.
$DA = \top$.

We leave it to the reader to show that $B_{KB*\alpha} \cup \{\beta\} = Th(\{p \wedge \neg q \wedge \neg r \wedge \neg s\})$ and $B_{KB*(\alpha\wedge\beta)} = Th(\{(\neg p \wedge \neg q \wedge r \wedge s) \vee (p \wedge \neg q \wedge \neg r \wedge \neg s)\})$.

■

VI. Examples

In this section we illustrate our approach by considering a number of examples.

Example 5 (Continuation of Example 4): Reconsider $KB = \langle\{p\}, \{p \Rightarrow q\}, \{\top\}\rangle$ and $\alpha = \neg p$. We showed that the revised knowledge base, $KB * \alpha$, is $\langle\{\neg p\}, \{p \Rightarrow q\}, \{\top\}\rangle$ and $B_{KB*\alpha} = Th(\{\neg p\})$. Suppose now that we learn that p holds after all. That is, $KB * \alpha$ should be revised by p. It is easily computed that the set of new observations is $\{p\}$ and so the new knowledge base is $\langle\{p\}, \{p \Rightarrow q\}, \{\top\}\rangle$. A belief set corresponding to this knowledge base is $Th(\{p, p \Rightarrow q\}) = Th(\{p, q\})$.

[12] $K \dotplus R$ is an abbreviation for $K \dotplus R \dot{-} \{\bot\}$.

[13] To fit our approach, some of these postulates have been slightly reformulated here.

[14] This counterexample has been given in [7] for a different revision operator. However it can also serve for our purposes.

To see that these results agree with our intuition, assume that p and q stand for "Tweety is a bird" and "Tweety flies", respectively. Our initial knowledge base consists of "Tweety is a bird" and "if Tweety is a bird, then she flies". This allows us to conclude that Tweety flies. When we learn that Tweety is not a bird, the conclusion that she flies must be withdrawn. But if we later learn that Tweety is a bird after all, this conclusion is again plausible.

■

Example 6: Consider now the knowledge base $KB = \langle \{b\}, \{b \Rightarrow f\}, \{p \Rightarrow \neg f, p \Rightarrow b\}\rangle$, where b, f and p are abbreviations for "Tweety is a bird", "Tweety flies" and "Tweety is a penguin", respectively. Suppose that the new observation, α, is p.

The new set of observations is $OB_1 = \{b, p\}$. (There is one EQ set, namely $\{b \Leftrightarrow b'\}$.) Thus, the revised knowledge base is $\langle \{b,p\}, \{b \Rightarrow f\}, \{p \Rightarrow \neg f, p \Rightarrow b\}\rangle$. A belief set corresponding to this knowledge base, $B_{KB*\alpha}$, is $\{b \Rightarrow f\} \dotplus \{b, p, p \Rightarrow \neg f, p \Rightarrow b\}$. Performing straightforward calculations, we get $B_{KB*\alpha} = Th(\{b, p, \neg f\})$.

Note that being a penguin is more specific than being a bird. So, the conclusion that Tweety does not fly is what intuition dictates.

■

Example 7: Consider the knowledge base $KB = \langle \{mo \vee jo, jc\}, \{\top\}, \{jc \Rightarrow \neg jo, mc \Rightarrow \neg mo\}\rangle$, where mo, mc, jo and jc stand for "John is in his office", "John is on the corridor", "Mary is in her office" and "Mary is on the corridor", respectively. Assume that a revision formula is $\alpha = \neg mo$.

A new set of observations is

$$OB_1 = \{mo \vee jo, jc\} \oplus \{\neg mo\} \ominus \{(jc \wedge jo) \vee (mc \wedge mo)\}.$$

We first look for maximal sets $EQ \subseteq \{mo \Leftrightarrow mo', jo \Leftrightarrow jo', jc \Leftrightarrow jc'\}$ such that $\{mo' \vee jo', jc'\} \cup \{\neg mo\} \cup EQ$ is consistent and $\{mo' \vee jo', jc'\} \cup \{\neg mo\} \cup EQ \nvdash (jc \wedge jo) \vee (mc \wedge mo)$. There are three such sets, namely $EQ_1 = \{jo \Leftrightarrow jo', mo \Leftrightarrow mo'\}$, $EQ_2 = \{jc \Leftrightarrow jc', mo \Leftrightarrow mo'\}$ and $EQ_3 = \{jo \Leftrightarrow jo', jc \Leftrightarrow jc'\}$. Accordingly, there are three FE formulae (see Algorithm 1), namely $FE_1 = \exists jc' \exists mo' \exists jo'.[(mo' \vee jo') \wedge jc' \wedge \neg mo \wedge (jo \Leftrightarrow jo') \wedge (mo \Leftrightarrow mo')$, $FE_2 = \exists jc' \exists mo' \exists jo'.[(mo' \vee jo') \wedge jc' \wedge \neg mo \wedge (jc \Leftrightarrow jc') \wedge (mo \Leftrightarrow mo')$ and $FE_3 = \exists jc' \exists mo' \exists jo'.[(mo' \vee jo') \wedge jc' \wedge \neg mo \wedge (jo \Leftrightarrow jo') \wedge (jc \Leftrightarrow jc')$. It is easily verified that FE_1, FE_2 and FE_3 reduce to $\neg mo \wedge jo$, $\neg mo \wedge jc$ and $\neg mo \wedge jc$, respectively. Thus, $OB_1 = \neg mo \wedge (jc \vee jo)$, and hence $KB * \alpha = \langle \{\neg mo \wedge (jc \vee jo)\}, \{\top\}, \{jc \Rightarrow \neg jo, mc \Rightarrow \neg mo\}\rangle$. A belief set corresponding to this knowledge base is given by $B_{KB*\alpha} = \{\top\} \dotplus \{\neg mo \wedge (jc \vee jo)\} \cup \{jc \Rightarrow \neg jo, mc \Rightarrow \neg mo\} = Th(\{\neg mo \wedge (jc \vee jo)\} \cup \{jc \Rightarrow \neg jo, mc \Rightarrow \neg mo\} = Th(\{\neg mo \wedge (jo \wedge \neg jc \vee jc \wedge \neg jo) \wedge (jc \Rightarrow \neg jo) \wedge (mo \Rightarrow \neg mc))$. Thus, after revision, we conclude that Mary is not in her office, whereas John is in his office or on the corridor, but not in both of these places.

■.

VII. PRIORITIZED BELIEF REVISION

It is not difficult to see that our revision operator specifies a priority ordering on different kinds of knowledge. The highest priority is assigned to domain axioms. The lower priority is connected with observations which are revised before defeasible statements. Finally, the lowest priority is assigned to defeasible statements. However, it turns out that sometimes it is necessary to provide additional priority ordering on defeasible statements. The following example will help to illustrate this point.[15]

Example 8: Consider the following facts.

Mary is an adult.
Typically, if Mary is an adult, then Mary is employed.
Typically, if Mary is a full-time student, then Mary is not employed.

This is represented by $KB = \langle OB, DS, DA\rangle$, where (here A, F and E stand for "Mary is an adult", "Mary is a full-time student" and "Mary is employed", respectively)

$OB = \{A\}$
$DS = \{A \Rightarrow E, \ F \Rightarrow \neg E\}$
$DA = \{\top\}$.

Suppose that we have learned that Mary is a full-time student. Therefore, KB should be revised by $\alpha = F$. It is immediately seen that $OB_1 = OB \oplus \{F\} \ominus \{\bot\} = OB \oplus \{F\} = \{A, F\}$. Thus the revised knowledge base, $KB * \alpha$, is $\langle \{A, F\}, \{A \Rightarrow E, \quad F \Rightarrow \neg E\}, \{\top\}\rangle$. To compute a set of beliefs corresponding to this knowledge base, we calculate $\{A \Rightarrow E, \quad F \Rightarrow \neg E\} \oplus \{A, F\}$. The set of atoms occurring both in DS and OB_1 is $\{A, F\}$. Thus we look for maximal $EQ \subseteq \{A \Leftrightarrow A', F \Leftrightarrow F'\}$ such that $\{A \Rightarrow E, \quad F \Rightarrow \neg E\} \cup \{A, F\} \cup EQ$ is consistent. There are two such sets, namely $EQ_1 = \{A \Leftrightarrow A'\}$ and $EQ_2 = \{F \Leftrightarrow F'\}$. Accordingly, there are two FE formulae (see Algorithm 1), namely $FE_1 = \exists A' \exists F'.(A' \Rightarrow E) \wedge (F' \Rightarrow \neg E) \wedge A \wedge F \wedge (A \Leftrightarrow A')$ which reduces to $A \wedge F \wedge E$ and $FE_2 = \exists A' \exists F'.(A' \Rightarrow E) \wedge (F' \Rightarrow \neg E) \wedge A \wedge F \wedge (F \Leftrightarrow F')$ which reduces to $A \wedge F \wedge \neg E$. Thus $B_{KB*\alpha} = Th(\{(A \wedge F \wedge E) \vee (A \wedge F \wedge \neg E)\} = Th(\{A, F\})$.

This result is intuitively unacceptable. Given that Mary is an adult full-time student, we should not be agnostic about her employment status. Rather, we should infer that Mary is unemployed.

There is an obvious conflict between the both defeasible statements. Given that Mary is an adult, there is a reason to believe that she is employed. Given that she is a full-time student, there is reason to believe the contrary. In consequence, neither of these conclusions can be inferred.

To resolve this problem, the priority assigned to $F \Rightarrow \neg E$ should be higher than that assigned to $A \Rightarrow E$. This simply means that the former of these formulae should be considered before the latter during the revision process. To achieve this effect, we divide the set DS into $DS_1 = \{A \Rightarrow E\}$ and

[15] This is a propositional version of a classical example presented first in [15] in the context of default reasoning (see also [16]).

$DS_2 = \{F \Rightarrow \neg E\}$ and define $B_{KB*\alpha}$ as $DS_1 \dotplus (DS_2 \oplus (OB_1 \cup DA))$ instead of $DS \dotplus (OB_1 \cup DA)$.

Recall that $OB_1 \cup DA = \{A, F\}$. Thus, $DS_2 \oplus (OB_1 \cup DA) = DS_2 \oplus \{A, F\} = \{A, F, \neg E\}$. Therefore, $B_{KB*\alpha} = \{DS_1 \dotplus \{A, F, \neg E\} = \{A \Rightarrow E\} \dotplus \{A, F, \neg E\} = Th(\{A, F, \neg E\})$.

This agrees with our intuition.

■

We are now ready to provide a formal definition of prioritized version of our revision operator. Consider a knowledge base $KB = \langle OB, DS, DA\rangle$. We begin by breaking DS into disjoint sets $DS_1, \ldots, DS_n$. The intention is that the elements of DS_n are given the highest priority, those of DS_{n-1} the second priority, etc. We express this by writing $DS_1 < DS_2 < \cdots < DS_n$.

Definition 7: Let $KB = \langle OB, DS, DA\rangle$ be a knowledge base and α be a revision formula. Let $OB_1 = OB \oplus \{\alpha\} \ominus \{\neg DA^\wedge\}$. The *prioritized belief revision of* KB *by* α with respect to priorities $DS_1 < DS_2 < \cdots < DS_n$, written $KB *^{[DS_1 < DS_2 < \cdots < DS_n]} \alpha$, is the formula

$$DS_1 \dotplus (DS_2 \oplus \cdots (DS_n \oplus (OB_1 \cup DA) \cdots).$$

■

It should be noted that the prioritized version of our belief operator does not change the revised knowledge base $K * \alpha$ which remains as in Definition 5.

VIII. Absorbing knowledge

So far we have assumed that a revision formula is an observation. In this section we redefine our revision operator under the assumption that this formula is a piece of knowledge.

Suppose first that the revising formula α represents a domain axiom. In this case we proceed as follows.

Definition 8: Let $KB = \langle OB, DS, DA\rangle$ and suppose that α is a revision formula representing a domain axiom. The revised knowledge base is defined by $KB * \alpha = \langle OB_1, DS, DA \cup \{\alpha\}\rangle$, where $OB_1 = OB \oplus \{\top\} \ominus \{\neg(DA \cup \{\alpha\})^\wedge\}$. Here $OB \oplus \{\top\} \ominus \{\neg(DA \cup \{\alpha\})^\wedge\}$ is a finite representation of the modified knowledge base corresponding to belief change scenario $\langle OB, \{\top\}, \{\neg(DA \cup \{\alpha\})^\wedge\}\rangle$. (See section III.)

■

Note that the addition of a new domain axiom may contradict a set of observations. This is the reason why this set is modified.

Definition 8 assumes that a new domain axiom is consistent with the set of the current domain axioms. This assumption seems natural. However, if one wanted to reject it, $KB * \alpha$ should be defined as $\langle OB_1, DS, DA \oplus \{\alpha\}\rangle$, where OB_1 is specified as before.

Example 9: This is a variant of Example 6. Let $KB = \langle\{p\}, \{b \Rightarrow f\}, \{p \Rightarrow b\}\rangle$, where b, f and p stand for "Tweety is a bird", "Tweety flies" and "Tweety is a penguin", respectively. Since there is no information that penguins do not fly, KB can be used to conclude that Tweety is a flier. This is indeed the case. We leave it to the reader to show that $f \in B_{KB}$.

Suppose now that we revise our knowledge base by a domain axiom $\alpha = (p \Rightarrow \neg f)$. $KB * \alpha = \langle OB_1, \{b \Rightarrow f\}, \{p \Rightarrow b, p \Rightarrow \neg f\}\rangle$, where $OB_1 = \{b\} \oplus \{\top\} \ominus \{\neg[(p \Rightarrow b \wedge (p \Rightarrow \neg f)]\} = \{b, p\}$. Thus $KB * \alpha = \langle\{b, p\}, \{b \Rightarrow f\}, \{p \Rightarrow b, p \Rightarrow \neg f\}\rangle$. Performing straightforward calculations, we get $\neg f \in B_{KB*\alpha}$. ■

Example 10: Let $KB = \langle\{p, q\}, \{\top\}, \{\top\}\rangle$ and suppose that a revision formula is a new domain axiom $\alpha = (p \Rightarrow \neg q)$.[16] $OB_1 = \{p, q\} \oplus \{\top\} \ominus \{\neg(\top \wedge (p \Rightarrow \neg q))\} = \{p, q\} \oplus \{\top\} \ominus \{p \wedge q\}$. It is easily verified that $OB_1 = \{p \vee q\}$ and $B_{KB*\alpha} = Th(\{p \Leftrightarrow \neg q\}$.

■

Suppose now that a revision formula α is a defeasible statement. In this case we proceed as follows.

Definition 9: Let $KB = \langle OB, DS, DA\rangle$ and suppose that α is a revision formula representing a defeasible statement. The revised knowledge base is defined by $KB * \alpha = \langle OB, DS_1, DA\rangle$, where $DS_1 = DS \oplus \{\alpha\}$.

■

Note that a belief set associated with $\langle OB, DS_1, DA\rangle$ is $DS_1 \dotplus (OB \cup DA)$. Therefore, there is no need to require consistency of $DS_1 \cup OB \cup DA$.

IX. Related work

Our formalization of belief revision is based on the assumption that observations should be separated from knowledge. Similar assumption is also made in [10]. However there are three important differences between both of the approaches.

(i) In [10] a knowledge base is a pair $\langle OB, A\rangle$, where OB is a set of observations and A is a set of domain axioms. Thus, formulae representing defeasible knowledge like "A bird typically flies" must be treated as observations. Example 1 shows that this may lead to strange results.

(ii) In [10] it is assumed that a new observation describes exactly what a reasoning agent knows at the moment about the aspect of the world the observation concerns. For instance, if all we believe is p and new observation is $p \vee q$, then after revision all we believe is $p \vee q$. This point will be discussed in detail in the next section.

(iii) In contrast to the belief revision operator presented here, the operator specified in [10] is based on eliminating non-redundant atoms.

X. Concluding remarks

The AGM postulates of belief revision introduce mathematically consistent and elegant theory. However, to obtain a clear, simple and manageable model, substantial simplifications have been made (see [6] for a discussion of this point). These simplifications often lead to unintuitive results.

In this paper we have rejected one of those simplifications, namely the assumption that a knowledge base is unstructured.

[16] Observe that α is inconsistent with the current set of observations.

Applying our revision operator to structured knowledge bases, allows us to express knowledge in a more flexible way. In particular, we can deal with scenarios that are difficult, or impossible, to model using existing approaches to belief revision. For instance, we can assign priorities to defeasible statements.

We distinguish between formulae representing observations and formulae representing knowledge. In addition, the latter formulae are divided into two groups. A natural question arises about the nature of various types of information represented by those formulae and agents' degree of freedom. Having observed a contradictory fact, should an agent be allowed to change domain axiom into a defeasible statement or completely reject new information? Consider the following example.[17]

Suppose that we know that all swans are white. What should we do when a black swan is observed? If we are sure that an observed object is a black swan, common sense suggests the change of the domain axiom stating that all swans are white into a defeasible statement stating that a typical swan is white. It is worth noting that this natural solution cannot be achieved if a knowledge base is unstructured.

Observe that neither observations nor domain axioms have influence on a set of defeasible statements in a revised knowledge base. Of course, it may happen that a defeasible statement is defeated by an observation or a domain axiom.[18] However, even in this case we keep such defeated defeasible statement in the revised knowledge base. To see why, consider the following example.

Let $KB = \langle\{\top\}, \{\neg g, \}, \{e \wedge \neg a \Rightarrow g\}\rangle$ where g, e and a stand for "John is guilty", "There is evidence that John is guilty" and "John has an alibi", respectively. A belief set corresponding to KB is $Th(\{\neg g \wedge (\neg e \vee a)\})$. During an investigation an evidence that John is guilty and has no alibi was found, i.e, KB should be revised by $\alpha = e \wedge \neg a$. It is easily seen that the modified knowledge base, $KB * \alpha$, is $\langle\{e \wedge \neg a\}, \{\neg g\}, \{e \wedge \neg a \Rightarrow g\}\rangle$ and a belief set corresponding to this knowledge base is $Th(\{e, \neg a, g\})$. Now, if we allow to override defeasible statements by observations or domain axioms, $\neg g$ should be withdrawn from a set of defeasible statements of $KB * \alpha$. But in this case, if we later learn that John has an alibi after all, we shall not be able to infer that he is innocent.

One of the AGM postulates states that whenever a revised formula α is consistent with a knowledge base KB, then $KB * \alpha$ is just $KB \cup \{\alpha\}$. In particular, if a revision formula is a consequence of a knowledge base, then the process of revision does not change the base.[19] It turns out that the plausibility of this postulate depends on the context. Two simple examples will help to illustrate this.

[17] This example is from [5].

[18] For this reason a belief set associated with a knowledge base $\langle OB, DS, DA\rangle$ is defined as $DS \dotplus (OB \cup DA)$ instead of $Th(OB \cup DS \cup DA)$.

[19] It seems that an operator defined in [10] is the only exception of this rule.

Assume first that during the day I observed that my neighbour's new car is green. At rainy night I saw the car again. This time I observed that the car's color is green or brown. Obviously, there is no need to change my belief set. It is plausible to assume that the car is green.

The second example concerns a weather forecast. Watching TV yesterday, I had learned that two days later it would rain in Warsaw. Watching TV today, I have learned that tomorrow it will rain or snow in Warsaw. Clearly, intuition dictates that today's forecast is more reliable, so that I should change my belief set.

From the above discussion we can see that the process of belief revision is extremely complex. Many factors have to be taken into consideration. There are still a large number of important open problems in this field of research which require inputs from both logic and philosophy for their solution [6].

REFERENCES

[1] C. Alchourron, P. Gardenfors and D. Makinson. On the Logic of Theory Change: Partial Meet Functions for Contraction and Revision. *Journal of Symbolic Logic*, 50:510-530, 1985.

[2] F. M. Brown. *Boolean Reasoning*. Kluwer Academic Publishers, 1990.

[3] J. Delgrande, T. Schaub. A Consistency-Based Approach for Belief Change. *Artificial Intelligence Journal*, 151, 1-2, 1-41, 2003.

[4] P. Doherty, W. Łukaszewicz, E. Madalińska-Bugaj. The PMA and Relativizing Minimal Change for Action Update. *Fundamenta Informaticae*, 2000, 44(1-2), 95-131.

[5] P. Gärdenfors. Belief revision: An Introduction. In: P. Gärdenfors (ed.) *Belief revision*, volume 29 of *Cambridge Tracts in Theoretical Computer Science*, Cambridge University Press, 1992, 1-28. Electronic version: http://www.cs.sfu.ca/CourseCentral/823/jim/papers/gardenfors.pdf.

[6] S.O. Hansson. Ten Philosophical Problems in Belief Revision. *Journal of Logic and Computation*, 2003, 18(1), 37-49

[7] H. Katsuno, A. O. Mendelzon. A Unified View of Propositional Knowledge Base Updates. *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, 1989, 1413-1419.

[8] H. Katsuno, A. O. Mendelzon. On the Difference between Updating a knowledge base and revising it. In P. G rdenfors, editor, *Belief Revision*. Cambridge University Press, 1992, 183-203.

[9] F. Lin, R. Reiter. Forget It! *Proceedings of the AAAI Fall Symposium on Relevance*, 1994, 154-159. Electronic version: http://www.cs.toronto.edu/kr/papers/forgetting.pdf

[10] E. Madalińska-Bugaj, W. Łukaszewicz. Belief Revision Revisited. *Proceedings of the 4th Mexican International Conference on Artificial Intelligence*, Lecture Notes in Computer Science, Springer, 2005, 31-40.

[11] E. Madalińska-Bugaj, W. Łukaszewicz. First-order Generalization of the MPMA Belief Update Operator. *Fundamenta Informaticae*, 94, 49-61, 2009.

[12] Y. Moses, Y. Shoham,. Belief as Defeasible Knowledge. *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, 1989, 1168-1173. (Extended version in *Artificial Intelligence Journal*, 64(2), 1993, 299-321.)

[13] P. Peppas. Belief Revision. In: F. van Harmelen, V. Lifschitz and B. Porter (eds.), *Handbook of Knowledge Representation*, Elsevier, 2008.

[14] D. Perlis. On the Consistency of Common-Sense Reasoning. /em Computational Intelligence, 2, 1987, 180-190.

[15] R. Reiter, G. Criscuolo. On Interacting Defaults. *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, 1981, 270-276.

[16] R. Reiter, G. Criscuolo. Some Representational Issues in Default Reasoning. *International Journal of Computers and Mathematics*, 9, 1983, 1-9.