

Prime Implicate-based Belief Revision Operators ¹

Meghyn Bienvenu² and Andreas Herzig³ and Guilin Qi⁴

1 INTRODUCTION

A belief revision operator can be seen as a function which takes as input a set of beliefs K and an input formula φ and outputs a new set of beliefs $K \star \varphi$. Many of the belief revision schemes that have been defined in the literature require additional input. The extra information they need comes in various forms: relations over subsets of the sets of beliefs [2], epistemic entrenchment relations [1], system of spheres [6], faithful orderings [7], etc.

In many applications we do not have such background information, which is why there is a need for revision operators which give good results without it. Unfortunately, the above approaches appear ill-suited to cases where we do not have any information regarding the relative importance of different beliefs or models. For example, if we accord equal importance to each of the beliefs (or each model of the beliefs or non-beliefs), which seems the most reasonable thing to do if we have no preference information, then these approaches result in the infamous *drastic revision operator* which gives up all old beliefs whenever the incoming information contradicts them.

All of the above belief revision schemes are insensitive to syntax: logically equivalent sets of beliefs are revised in the same way, and logically equivalent input formulas lead to the same result. The so-called formula-based approaches, like the full meet [5, 4] and cardinality-maximizing base revision operators [5, 9], abandon the postulate of insensitivity to syntax, and allow e.g. the set of beliefs $K_1 = \{a, b\}$ to be revised differently from $K_2 = \{a \wedge b\}$. Such approaches can do without extra information: they do not collapse into the drastic revision operator.

There are only very few belief revision operators that are both insensitive to syntax and independent of extra information. The most prominent one is Dalal's [3]. It is often called model-based: revision is identified with a move from the models of K to those models of φ that are closest in terms of the Hamming distance. Two other model-based revision operators exist: Weber's [12] and Satoh's [11].

In this paper we propose two revision operators which are formula-based yet syntax-insensitive, and do not rely on background information. Our operators are obtained by first replacing the belief base by its set of prime implicates and then applying either the full meet or the cardinality-maximizing base revision operators. The prime implicates of a belief base, defined as its logically strongest clausal consequences, can be seen as the primitive semantic components of the belief base, from which all other beliefs can be derived. We argue that when no extra information is available, prime implicates provide a natural and interesting way of representing a set of beliefs. Moreover, the fact that equivalent sets of formulae have the same sets of

prime implicates guarantees the syntax-insensitivity of our operators.

2 FORMAL PRELIMINARIES

We consider a propositional language built out of a finite set of atoms and the usual Boolean connectives. We suppose the latter includes the 0-ary connective \perp . We will use $\mathcal{V}(\varphi)$ to refer to the set of atoms occurring in φ . A *belief base* is a finite set of propositional formulae. Where convenient, we will identify a belief base with the conjunction of its elements. We will use $\bigvee K$ to denote the disjunction of the elements in the belief base K . A *literal* is either an atom or the negation of an atom, and a *clause* is a disjunction of literals.

Prime implicates (cf. [8]) are defined as the logically strongest clausal consequences of a formula. By definition, if π is a prime implicate of φ , then so too are all clauses equivalent to π . To simplify the presentation, we will choose a representative for each equivalence class of clauses, and we let $\Pi(\varphi)$ denote the set of representatives of equivalence classes of prime implicates of φ .

We define the *minimal language* of a formula φ , written $\mathcal{V}_0(\varphi)$, to be the set of atoms occurring in every formula φ' which is equivalent to φ . A set $\{A_1, \dots, A_n\}$ of sets of atoms is a *splitting* of a belief base K if and only if the A_i partition $\mathcal{V}_0(K)$ and there exist formulae $\varphi_1, \dots, \varphi_n$ such that $K \equiv \bigwedge_{i=1}^n \varphi_i$ and $\mathcal{V}(\varphi_i) \subseteq A_i$ for all i . A splitting $\{A_1, \dots, A_n\}$ of K is a *finest splitting* of K just in the case that if $\{A'_1, \dots, A'_n\}$ is another splitting of K , then for every A_i there is some A'_j such that $A_i \subseteq A'_j$. It was shown in [10] that every belief base has a unique finest splitting.

We will use $K \perp \varphi$ and $K \perp_{Card} \varphi$ to denote respectively the set of inclusion- and cardinality-maximal subsets of K consistent with $\neg\varphi$.

3 PROPOSED REVISION OPERATORS

Our first revision operator \star_{Π} conjoins the input φ and the disjunction of the maximal subsets of $\Pi(K)$ consistent with φ . It is essentially the same as the syntactic full meet base revision operator [5, 4] except that instead of dealing directly with the formulae in the belief base we deal with the prime implicates of the belief base.

Definition 1. Let K be a belief base and φ be a formula. Then the *prime implicate-based full meet revision operator*, written \star_{Π} , is defined as follows:

$$K \star_{\Pi} \varphi = \varphi \wedge \bigvee (\Pi(K) \perp \neg\varphi)$$

We illustrate the functioning of our operator on some examples:

Example 2. Let $K = \{a \vee b, a \vee c\}$ and $\varphi = \neg a \wedge \neg b$. We have $\Pi(K) = K$, and $\Pi(K) \perp \neg\varphi = \{\{a \vee c\}\}$, so the result of revising K by φ is $\neg a \wedge \neg b \wedge (a \vee c) \equiv \neg a \wedge \neg b \wedge c$.

¹ The third author acknowledges partial support by the EU under the IST project NeOn (IST-2006-027595, <http://www.neon-project.org/>)

² IRIT-Université Paul Sabatier, France, bienvenu@irit.fr

³ IRIT-CNRS, France, herzig@irit.fr

⁴ Institute AIFB, Universität Karlsruhe, Germany, gqi@aifb.uni-karlsruhe.de

Example 3. Let $K = \{a \vee c, \neg b \vee d, \neg a \vee b\}$ and let $\varphi = \neg c \wedge \neg d$. Then $\Pi(K) = \{a \vee c, \neg b \vee d, \neg a \vee b, b \vee c, \neg a \vee d, c \vee d\}$. The maximal subsets of $\Pi(K)$ consistent with φ are $P_1 = \{a \vee c, \neg b \vee d\}$, $P_2 = \{a \vee c, \neg a \vee b, b \vee c\}$, $P_3 = \{\neg b \vee d, \neg a \vee b, \neg a \vee d\}$, and $P_4 = \{\neg a \vee b, b \vee c, \neg a \vee d\}$. Now $P_1 \wedge \neg c \wedge \neg d \equiv a \wedge \neg b \wedge \neg c \wedge \neg d$, $P_2 \wedge \neg c \wedge \neg d \equiv a \wedge b \wedge \neg c \wedge \neg d$, $P_3 \wedge \neg c \wedge \neg d \equiv \neg a \wedge \neg b \wedge \neg c \wedge \neg d$, and $P_4 \wedge \neg c \wedge \neg d \equiv \neg a \wedge b \wedge \neg c \wedge \neg d$, so $K \star_{\Pi} \varphi \equiv \neg c \wedge \neg d$.

In the last example, none of the prime implicates from K can be inferred from the revised base $K \star_{\Pi} \varphi$. This is because our operator takes the disjunction of *all* the inclusion-maximal subsets consistent with the revision formula, which means that those prime implicates which do not appear in every inclusion-maximal subset can be lost when we take the disjunction.

The solution lies in selecting only some of the inclusion-maximal subsets. If we have no information regarding the importance of different beliefs, as we assume here, there is no sure way of choosing among the subsets. One reasonable heuristic is to accord equal importance to each of the prime implicates, and hence to prefer those subsets which contain the most prime implicates. This leads us to propose a second revision operator which selects only those *cardinality-maximal* subsets consistent with the revision formula.

Definition 4. Let K be a belief base and φ be a formula. Then the *prime implicate-based cardinality-maximizing revision operator*, written $\star_{\Pi, Card}$, is defined as follows:

$$K \star_{\Pi, Card} \varphi = \varphi \wedge \bigvee (\Pi(K) \perp_{Card} \neg \varphi)$$

The operator $\star_{\Pi, Card}$ can be seen as a syntax-insensitive version of the cardinality-maximizing base revision operator [5, 9].

Example 5. Let K and φ be as in Example 3. P_2 , P_3 , and P_4 are the cardinality-maximal subsets that are consistent with φ . So we have $K \star_{\Pi, Card} \varphi \equiv (\neg a \vee b) \wedge \neg c \wedge \neg d$, which is logically stronger than $\neg c \wedge \neg d$ which is obtained using \star_{Π} .

3.1 Properties of Our Operators

Revision operators are often judged based on whether they satisfy the well-known AGM postulates [2]. These postulates are formulated for logically closed sets of formulae (belief sets), but they can be modified so as to apply to belief bases. The modified postulates (omitted for lack of space) are known as the KM postulates [7].

Our first operator satisfies the first five KM postulates but fails to satisfy the last one.

Proposition 6. \star_{Π} satisfies **KM1-KM5**, but falsifies **KM6**.

This proposition is not surprising since Katsuno and Mendelzon showed in [7] that **KM6** ensures that the *faithful assignment* corresponding to the revision operator is a total pre-order.⁵ As our prime implicate-based full meet operator uses inclusion to compare subsets of prime implicates, it induces a partial and not a total pre-order over the set of interpretations.

Katsuno and Mendelzon argued however in [7] that requiring the faithful assignment to be total may be too strong in practice, and they proposed to replace **KM6** with weaker postulates **KM7** and **KM8**. Since they are less well-known, we recall them here:

KM7 If $K \star \varphi_1 \models \varphi_2$ and $K \star \varphi_2 \models \varphi_1$ then $K \star \varphi_1 \equiv K \star \varphi_2$.

⁵ A faithful assignment maps a belief base K to a pre-order \leq_K over the set of all interpretations of the language.

KM8 $(K \star \varphi_1) \wedge (K \star \varphi_2) \models K \star (\varphi_1 \vee \varphi_2)$.

We show that both of these postulates are satisfied by our operator.

Proposition 7. \star_{Π} satisfies **KM7** and **KM8**.

Our cardinality-based operator satisfies all KM postulates.

Proposition 8. $\star_{\Pi, Card}$ satisfies **KM1-KM6**.

The AGM/KM postulates have been criticized for admitting revision operators that discard beliefs that have no real connection with the incoming information. For instance, there are AGM/KM operators for which $(a \wedge b) \star \neg a \not\models b$, even though intuitively we expect b to survive the revision. In an attempt to remedy this, Parikh [10] proposed an additional postulate which can be formulated as follows:

Relevance If K is satisfiable and $K \models \varphi$ and $K \star \psi \not\models \varphi$, then there is some set of atoms A in the finest splitting of K such that both $\mathcal{V}_0(\varphi) \cap A \neq \emptyset$ and $\mathcal{V}_0(\psi) \cap A \neq \emptyset$.

We can show that our revision operators satisfy this postulate:

Proposition 9. \star_{Π} and $\star_{\Pi, Card}$ satisfy **Relevance**.

3.2 Comparison With Other Operators

The following proposition concerns the relation between our operators and the model-based operators mentioned in the introduction.

Proposition 10.

1. Our operators sometimes yield logically stronger revised bases than the Dalal, Weber, and Satoh operators.
2. Our revision operators sometimes yield logically weaker revised bases than the Dalal and Satoh operators.

Proof. For (1), consider Example 2. For (2), consider Example 3.

REFERENCES

- [1] *Knowledge in Flux: Modeling the Dynamics of Epistemic States*, MIT Press, 1988.
- [2] C. Alchourron, P. Gärdenfors, and D. Makinson, ‘On the logic of theory change: Partial meet contraction and revision functions’, *Journal of Symbolic Logic*, **50**(2), 510–530, (1985).
- [3] M. Dalal, ‘Investigations into a theory of knowledge base revision’, in *Proceedings of the Seventh National Conference on Artificial Intelligence*, (475–479).
- [4] R. Fagin, J. Ullman, and M. Vardi, ‘On the semantics of updates in databases’, in *Proceedings of the Second ACM SIGACT-SIGMOD Symposium on Principles of Database Systems (PODS-83)*, pp. 352–365, (1983).
- [5] Matthew L. Ginsberg, ‘Counterfactuals’, *Artificial Intelligence*, **30**(1), 35–79, (1986).
- [6] A. Grove, ‘Two modelings for theory change’, *Journal of Philosophical Logic*, (1988).
- [7] H. Katsuno and A. Mendelzon, ‘Propositional knowledge base revision and minimal change’, *Artificial Intelligence*, **52**(3), 263–294, (1991).
- [8] P. Marquis, *Handbook on Defeasible Reasoning and Uncertainty Management Systems*, volume 5, chapter Consequence Finding Algorithms, 41–145, Kluwer, 2000.
- [9] B. Nebel, *Handbook of Defeasible Reasoning and Uncertainty Management Systems, Volume 3: Belief Change*, chapter How Hard is it to Revise a Belief Base?, Kluwer, 1998.
- [10] R. Parikh, *Logic, Language, and Computation*, volume 2, chapter Beliefs, belief revision, and splitting languages, CSLI Publications, 1999.
- [11] K. Satoh, ‘Nonmonotonic reasoning by minimal belief revision’, in *Proceedings of the Fifth International Conference Generation Computer Systems (FGCS-88)*, pp. 455–462, (1988).
- [12] A. Weber, ‘Updating propositional formulas’, in *Proceedings of the First Conference on Expert Database Systems*, pp. 487–500, (1986).