

## On Computational Aspects of Iterated Belief Change

Nicolas Schwind<sup>1\*</sup>, Sébastien Konieczny<sup>2</sup>, Jean-Marie Lagniez<sup>3</sup> and Pierre Marquis<sup>2,4</sup>

<sup>1</sup>National Institute of Advanced Industrial Science and Technology, Tokyo, Japan

<sup>2</sup>CRIL - CNRS, Université d'Artois, Lens, France

<sup>3</sup>Huawei Technologies Ltd, Boulogne-Billancourt, France

<sup>4</sup>Institut Universitaire de France

nicolas-schwind@aist.go.jp, konieczny@cril.fr, jmlagniez@gmail.com, marquis@cril.fr

### Abstract

Iterated belief change aims to determine how the belief state of a rational agent evolves given a sequence of change formulae. Several families of iterated belief change operators (revision operators, improvement operators) have been pointed out so far, and characterized from an axiomatic point of view. This paper focuses on the inference problem for iterated belief change, when belief states are represented as a special kind of stratified belief bases. The computational complexity of the inference problem is identified and shown to be identical for all revision operators satisfying Darwiche and Pearl's (R\*1-R\*6) postulates. In addition, some complexity bounds for the inference problem are provided for the family of soft improvement operators. We also show that a revised belief state can be computed in a reasonable time for large-sized instances using SAT-based algorithms, and we report empirical results showing the feasibility of iterated belief change for bases of significant sizes.

### 1 Introduction

Belief revision theory aims to study how to incorporate in the beliefs of an agent a new piece of information that (typically) contradicts them [Alchourrón *et al.*, 1985; Gärdenfors, 1988; Hansson, 1999; Fermé and Hansson, 2011]. Whereas the standard belief revision literature has been focused on the single-step change process, considering the iterated case, i.e., to determine what happens when several successive revision steps occur, is an important issue, that gave rise to a significant amount of work where iterated change operators have been defined and investigated [Lehmann, 1995; Darwiche and Pearl, 1997; Nayak *et al.*, 2003; Booth and Meyer, 2006; Rott, 2006; Jin and Thielscher, 2007; Konieczny and Pino Pérez, 2008; Konieczny *et al.*, 2010].

Improvement operators have been proposed as a general class of iterated change operators [Konieczny and Pino Pérez, 2008; Konieczny *et al.*, 2010]. In this paper, we will study two interesting subclasses of improvement operators from a computational perspective. The first one is

the class of iterated revision operators, which require the change formula to be entailed in the revised set of beliefs of the agent [Darwiche and Pearl, 1997; Nayak *et al.*, 2003; Booth and Meyer, 2006]. The second one is the class of soft improvement operators: in contrast to iterated revision operators, these operators aim to increase in a “minimal way” the plausibility of the change formula in the beliefs of the agent [Konieczny and Pino Pérez, 2008; Konieczny *et al.*, 2010].

In iterated belief change, it is standard to assume that the current set of beliefs of an agent is given by an epistemic state, i.e., an abstract object  $\Phi$  from which the actual beliefs of the agent, denoted by  $B(\Phi)$ , can be extracted. In this paper, the computational complexity of iterated belief change operators  $\star$  is investigated when epistemic states are represented by compact world rankings (CWRs). This representation will be formally introduced later. For the moment let us just roughly describe a CWR as a special kind of stratified belief base (an ordered set of formulae) that encodes a total preorder over the interpretations. We focus on the (one-step) inference problem: given a CWR  $\Phi = (\varphi_1, \dots, \varphi_n)$  and two formulae  $\mu, \alpha$ , is it the case that  $B(\Phi \star \mu) \models \alpha$ ? The computational complexity of this problem is identified for several families of iterated change operators  $\star$ , and shown as “mildly” hard (i.e., solvable using a constant number or a logarithmic number of calls to an NP oracle). More in detail, we show that:

1. for any iterated revision operator satisfying postulates **(R\*1)-(R\*6)** [Darwiche and Pearl, 1997] the inference problem is in  $\Theta_2^P$ , i.e., only a logarithmic number of calls to an NP oracle is necessary in the general case to decide whether a formula is entailed by a revised epistemic state;
2. for the three soft improvement operators pointed out in the literature [Konieczny *et al.*, 2010], the inference problem is shown to be in the Boolean Hierarchy **BH** [Wagner, 1987; Cai *et al.*, 1988], i.e., only a bounded number of calls to an NP oracle is needed to solve the problem; the exact complexity of the problem within **BH** (completeness results) is reported for each of those three operators;
3. hardness results for the inference problem, that rely only on postulates and apply to all operators satisfying them, have been identified as well. Thus, for any iterated revision operator satisfying the postulates **(R\*1)-(R\*6)**, the inference problem is shown  $\Theta_2^P$ -hard (hence, given the re-

\*Contact Author

sults in 1., it is  $\Theta_2^P$ -complete). For any soft improvement operator, it is shown coNP(3)-hard.

Given the performance of modern SAT solvers (that are NP oracles), these results suggest that the existence of inference algorithms for iterated belief change that prove efficient enough to tackle large-sized instances is not utopic. In order to verify it, we implemented SAT-based algorithms for the iterated belief change of CWRs (seven operators have been considered). More precisely, for each change operator under consideration, we implemented the change operation  $\Phi \star \mu$ , i.e., we designed a procedure for deriving a CWR representing  $\Phi \star \mu$  from a CWR representing  $\Phi$ , and from  $\mu$ . Interestingly, with epistemic states  $\Phi$  represented as CWRs, the representation of  $\Phi \star \mu$  does not exponentially blow up in terms of size: the number of strata in the corresponding CWR representation of  $\Phi \star \mu$  is at most twice the one of  $\Phi$ . As an interesting consequence, building the CWR  $\Phi \star \mu$  can be done in polynomial time given an NP oracle. We evaluated our SAT-based algorithms for iterated belief change on many benchmarks. The empirical results we have obtained confirm the feasibility of iterated belief change from the practical side.

The proofs of propositions can be found on <http://www.cril.fr/~koniczny/SKLM-IJCAI20-long.pdf>

## 2 Preliminaries on Iterated Belief Change

Let  $\mathcal{L}_{\mathcal{P}}$  be a propositional language built up from a finite set of propositional variables  $\mathcal{P}$  and the usual connectives.  $\perp$  (resp.  $\top$ ) is the Boolean constant always false (resp. true). An interpretation (or world) is a mapping from  $\mathcal{P}$  to  $\{0, 1\}$ .  $\models$  denotes logical entailment,  $\equiv$  logical equivalence, and  $[\varphi]$  denotes the set of models of the formula  $\varphi$ . In iterated change, it is standard to assume that the set of beliefs of an agent is represented by an *epistemic state* (ES for short)  $\Phi$ , which represents (i) the actual beliefs of the agent, i.e., a propositional formula denoted by  $B(\Phi)$ , and (ii) some conditional information guiding the revision process.

For conventional purposes, in the rest of the paper the symbols  $\circ$ ,  $\bullet$  and  $\star$  will respectively denote an iterated revision operator, a soft improvement operator, and any iterated belief change operator.

### 2.1 Iterated Revision

Let us start with Darwiche and Pearl's iterated revision:

**Definition 1 (DP-AGM revision operator)** [Darwiche and Pearl, 1997]. A DP-AGM revision operator  $\circ$  is an operator associating an ES  $\Phi$  and a formula  $\mu$  with a new ES  $\Phi \circ \mu$ , such that for each ES  $\Phi$  and all formulae  $\mu, \mu'$ :

- (R\*1)  $B(\Phi \circ \mu) \models \mu$ ;
- (R\*2) If  $B(\Phi) \wedge \mu \not\models \perp$ , then  $B(\Phi \circ \mu) \equiv B(\Phi) \wedge \mu$ ;
- (R\*3) If  $\mu \not\models \perp$ , then  $B(\Phi \circ \mu) \not\models \perp$ ;
- (R\*4) If  $\mu \equiv \mu'$ , then  $B(\Phi \circ \mu) \equiv B(\Phi \circ \mu')$ ;
- (R\*5)  $B(\Phi \circ \mu) \wedge \mu' \models B(\Phi \circ (\mu \wedge \mu'))$ ;
- (R\*6) If  $B(\Phi \circ \mu) \wedge \mu' \not\models \perp$ , then  $B(\Phi \circ (\mu \wedge \mu')) \models B(\Phi \circ \mu) \wedge \mu'$ .

Darwiche and Pearl also provided a characterization of DP-AGM operators in terms of total preorders over worlds:

**Definition 2 (faithful assignment).** A function  $\Phi \mapsto \preceq_{\Phi}$  that maps each ES  $\Phi$  to a total preorder<sup>1</sup> over worlds  $\preceq_{\Phi}$  is a faithful assignment iff for all worlds  $\omega, \omega'$ :

1. If  $\omega \models B(\Phi)$  and  $\omega' \models B(\Phi)$ , then  $\omega \simeq_{\Phi} \omega'$ ;
2. If  $\omega \models B(\Phi)$  and  $\omega' \not\models B(\Phi)$ , then  $\omega \prec_{\Phi} \omega'$ ;

**Proposition 1** ([Darwiche and Pearl, 1997]). An operator  $\circ$  is a DP-AGM operator iff there exists a faithful assignment  $\Phi \mapsto \preceq_{\Phi}$  such that for each ES  $\Phi$  and each formula  $\mu$ ,  $[B(\Phi \circ \mu)] = \min([\mu], \preceq_{\Phi})$ .

One can see from Prop. 1 that given an ES  $\Phi$  and a formula  $\mu$ , the set of models of  $B(\Phi \circ \mu)$  can be characterized independently of the choice of the DP-AGM revision operator  $\circ$ . However, postulates (R\*1-R\*6) impose no restriction on the rest of the ordering  $\preceq_{\Phi \circ \mu}$ . Thus four additional “rigidity” postulates (CR1-CR4) have been introduced [Darwiche and Pearl, 1997]. In addition, postulate (PR) was also introduced to also require a strict increase of the plausibility of models of  $\mu$  within  $\preceq_{\Phi}$  [Booth and Meyer, 2006; Jin and Thielscher, 2007]:<sup>2</sup>

- (CR1) If  $\omega \models \mu$  and  $\omega' \models \mu$ , then  $\omega \preceq_{\Phi} \omega' \Leftrightarrow \omega \preceq_{\Phi \circ \mu} \omega'$ ;
- (CR2) If  $\omega \not\models \mu$  and  $\omega' \not\models \mu$ , then  $\omega \preceq_{\Phi} \omega' \Leftrightarrow \omega \preceq_{\Phi \circ \mu} \omega'$ ;
- (CR3) If  $\omega \models \mu$  and  $\omega' \not\models \mu$ , then  $\omega \prec_{\Phi} \omega' \Rightarrow \omega \prec_{\Phi \circ \mu} \omega'$ ;
- (CR4) If  $\omega \models \mu$  and  $\omega' \not\models \mu$ , then  $\omega \preceq_{\Phi} \omega' \Rightarrow \omega \preceq_{\Phi \circ \mu} \omega'$ ;
- (PR) If  $\omega \models \mu$  and  $\omega' \not\models \mu$ , then  $\omega \preceq_{\Phi} \omega' \Rightarrow \omega \prec_{\Phi \circ \mu} \omega'$ .

DP-AGM operators satisfying (CR1-CR2) and (PR) are called *admissible* operators [Booth and Meyer, 2006]. Noteworthy, these operators also satisfy (CR3-CR4).

Let us now introduce four well-known DP-AGM operators from the literature. For space reasons, each one of them is semantically characterized by additional postulates on their corresponding faithful assignment. This also allows one to see how the revised ES  $\Phi \circ \mu$  is constructed in terms of total preorders (an illustrative example is given in Fig. 1).

**Boutilier's natural revision.** Denoted by  $\circ_B$ , it is the DP-AGM operator satisfying (CR1-CR4) and the following additional postulate [Boutilier, 1996]:

- (CBR) If  $\omega, \omega' \not\models B(\Phi \circ_B \mu)$ , then  $\omega \preceq_{\Phi} \omega' \Leftrightarrow \omega \preceq_{\Phi \circ_B \mu} \omega'$ .

This operator does not satisfy (PR).

**Nayak's lexicographic revision.** Denoted by  $\circ_N$ , it is the DP-AGM operator satisfying (CR1-CR4) and the following additional postulate [Nayak et al., 2003]:

- (R) If  $\omega \models \mu$  and  $\omega' \not\models \mu$ , then  $\omega \prec_{\Phi \circ_N \mu} \omega'$ .

This operator satisfies (PR).

<sup>1</sup>For each preorder  $\preceq$ ,  $\simeq$  denotes the corresponding indifference relation, and  $\prec$  the corresponding strict ordering.

<sup>2</sup>For space reasons, we only provide the semantic versions of the postulates, i.e., in terms of properties on faithful assignments.

**Lehmann's ranked revision.** Denoted by  $\circ_L$ , it is the DP-AGM operator characterized by a *widening ranked model* [Lehmann, 1995]. Although Lehmann initially characterized the operator in terms of functions mapping ordinals to non-empty subsets of worlds, it can equivalently be characterized in terms of a faithful assignment satisfying the following properties:

- (L1) If  $\omega, \omega' \not\models \mu$  and  $\forall \omega'' \models \mu, \omega \preceq_{\Phi} \omega''$ ,  
then  $\omega \preceq_{\Phi \circ_L \mu} \omega'$ ;
- (L2) If  $\exists \omega'' \models \mu$  s.t.  $\omega'' \prec_{\Phi} \omega'$ ,  
then  $\omega \preceq_{\Phi} \omega' \Leftrightarrow \omega \preceq_{\Phi \circ_L \mu} \omega'$ ;

This operator satisfies (CR1), (CR3) and (CR4) but it does not satisfy (CR2) or (PR).

**Booth and Meyer's restrained revision.** Denoted by  $\circ_{BM}$ , it is the DP-AGM operator satisfying (CR1-CR2), (PR), and the following additional postulate [Booth and Meyer, 2006]:

- (DR) If  $\omega \models \mu, \omega \not\models B(\Phi \circ_{BM} \mu)$  and  $\omega' \not\models \mu$ ,  
then  $\omega' \prec_{\Phi} \omega \Rightarrow \omega' \prec_{\Phi \circ_{BM} \mu} \omega$ .

## 2.2 Soft Improvement

Soft improvement operators, unlike DP-AGM ones, do not satisfy the success postulate (R\*1) [Konieczny *et al.*, 2010]. Instead, they satisfy the following one:

- (II) There exists  $n > 0$  such that  $B(\Phi \bullet^n \mu) \models \mu$ ,

where  $\Phi \bullet^n$  is inductively defined as  $\Phi \bullet^1 \alpha = \Phi \bullet \alpha$  and for each  $n \geq 1$ ,  $\Phi \bullet^{n+1} \alpha = (\Phi \bullet^n \alpha) \bullet \alpha$ .

**Definition 3** (soft improvement operator). A soft improvement operator is an operator  $\bullet$  associating an ES  $\Phi$  and a formula  $\mu$  with a new ES  $\Phi \bullet \mu$  such that postulates<sup>3</sup> (II-II0) from [Konieczny *et al.*, 2010] are satisfied.

Konieczny *et al.* [2010] also provided the following representation theorem :

**Definition 4** (soft gradual assignment). A strong faithful assignment is a faithful assignment such that for each  $n > 0$ , and each ES  $\Phi$ , we have  $\preceq_{\Phi \bullet \alpha_1 \bullet \dots \bullet \alpha_n} = \preceq_{\Phi \bullet \beta_1 \bullet \dots \bullet \beta_n}$  when all formulae  $\alpha_i, \beta_i, i \in \{1, \dots, n\}$  are such that  $\alpha_i \equiv \beta_i$ .

A soft gradual assignment is a strong faithful assignment satisfying (CR1), (CR2), (PR) and the following property:

- (S4) If  $\omega \models \mu$  and  $\omega' \not\models \mu$ , then  $\omega' \prec_{\Phi} \omega \Rightarrow \omega' \preceq_{\Phi \bullet \mu} \omega$ .

**Proposition 2** ([Konieczny *et al.*, 2010]). An operator  $\bullet$  is a soft improvement operator iff there exists a soft gradual assignment  $\Phi \mapsto \preceq_{\Phi}$  such that for each ES  $\Phi$  and each formula<sup>4</sup>  $\mu$ ,  $[B(\Phi \bullet \mu)] = \min([\mu], \preceq_{\Phi})$ .

As opposite to DP-AGM operators, when performing a change  $\Phi \bullet \mu$  soft improvement operators  $\bullet$  require the plausibility of all models of  $\mu$  w.r.t.  $\preceq_{\Phi}$  to be increased only to a "small" extent. This behavior is reflected by (S4).

Lastly, let us introduce the (semantic) characterization of the three soft improvement operators found in [Konieczny *et al.*, 2010].

<sup>3</sup>Postulates are not recalled here. For space reasons we will focus on the corresponding assignment properties. The reader is invited to check [Konieczny *et al.*, 2010] for more details.

<sup>4</sup> $\Phi \bullet \mu = \Phi \bullet^n \mu$  where  $n$  is the smallest integer such that  $B(\Phi \bullet^n \mu) \models \mu$ .

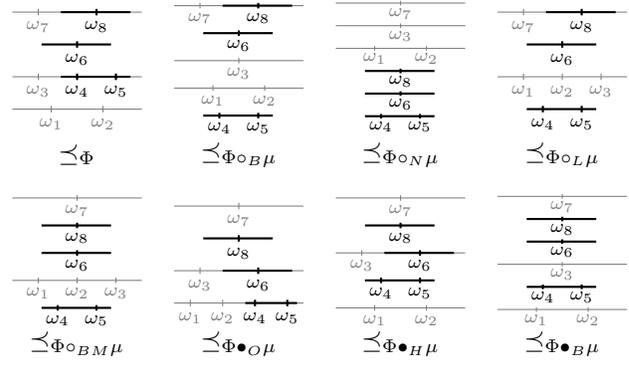


Figure 1: Illustration of an ES  $\Phi$  as a total preorder over worlds, and the changed ES  $\Phi \star \mu$  for  $\star \in \{\circ_B, \circ_N, \circ_L, \circ_{BM}, \bullet_O, \bullet_H, \bullet_B\}$  with  $[\mu] = \{\omega_4, \omega_5, \omega_6, \omega_8\}$ .

**One-improvement.** Denoted by  $\bullet_O$ , it is the soft improvement operator satisfying the following additional property. For all  $\omega, \omega', \omega \ll_{\Phi} \omega'$  is a shortcut for  $(\omega \prec \omega'$  and  $\nexists \omega'', \omega \prec \omega'' \prec \omega')$ :

- (S5) If  $\omega \models \mu$  and  $\omega' \not\models \mu$ , then  $\omega' \ll_{\Phi} \omega \Rightarrow \omega \preceq_{\Phi \bullet_O \mu} \omega'$ .

**Half-improvement.** Denoted by  $\bullet_H$ , it is the soft improvement operator satisfying the following two additional properties:

- (SH1) If  $\omega \models \mu, \omega' \not\models \mu, \omega' \ll_{\Phi} \omega$  and  $\nexists \omega'' \not\models \mu$  such that  $\omega'' \simeq_{\Phi} \omega$ , then  $\omega \preceq_{\Phi \bullet_H \mu} \omega'$ ;
- (SH2) If  $\omega \models \mu, \omega' \not\models \mu, \omega' \ll_{\Phi} \omega$  and  $\exists \omega'' \not\models \mu$  such that  $\omega'' \simeq_{\Phi} \omega$ , then  $\omega' \prec_{\Phi \bullet_H \mu} \omega$ .

**Best-improvement.** Denoted by  $\bullet_B$ , it is the soft improvement operator satisfying the following two additional properties. A formula  $\alpha$  is *separated* in  $\preceq$  for a given total preorder  $\preceq$  iff for all  $\omega, \omega'$ , if  $\omega \models \alpha$  and  $\omega' \not\models \alpha$  then  $\omega \not\prec \omega'$ :

- (SB1) If  $\omega \models \mu, \omega' \not\models \mu, \omega' \ll_{\Phi} \omega$  and  $\mu$  is separated in  $\preceq_{\Phi}$ , then  $\omega \preceq_{\Phi \bullet_B \mu} \omega'$ ;
- (SB2) If  $\omega \models \mu, \omega' \not\models \mu, \omega' \prec_{\Phi} \omega$  and  $\mu$  is not separated in  $\preceq_{\Phi}$ , then  $\omega' \prec_{\Phi \bullet_B \mu} \omega$ .

## 3 Complexity Results

From now on, we assume that each change operator  $\star$  is characterized in terms of a faithful assignment (cf. Def. 2 and 4), and each ES  $\Phi$  is represented by a *compact world ranking* (CWR for short):

**Definition 5** (compact world ranking). A compact world ranking (CWR) is a vector of consistent formulae  $(\varphi_1, \dots, \varphi_n)$ ,  $n \geq 1$ , such that for all  $i, j, i \neq j \Rightarrow \varphi_i \wedge \varphi_j \models \perp$  and  $\bigvee_{i=1}^n \varphi_i$  is valid.

Stated otherwise, a CWR is a total ordering over a set of consistent formulae whose sets of models are jointly exhaustive and pairwise disjoint (JEPD). To obtain the ES  $\Phi$  (in terms of faithful assignment) corresponding to a given CWR  $(\varphi_1, \dots, \varphi_n)$ , one simply sets  $\omega \preceq_{\Phi} \omega'$  iff  $\omega \models \varphi_i$  and  $\omega' \models \varphi_j$  for some  $i, j$  s.t.  $i < j$ . Doing so, it is easy to see that  $B(\Phi) = \varphi_1$ . Abusing notations, in the following  $\Phi$

will denote both an ES and the CWR representing it. Lastly, in a CWR  $\Phi = (\varphi_1, \dots, \varphi_n)$ , for each  $i \in \{1, \dots, n\}$ , the formula  $\varphi_i$  is also called the  $i^{\text{th}}$  stratum of  $\Phi$ , and  $\Phi$  is said to have  $n$  strata.

A CWR can be seen as a special kind of stratified belief base (SBB), as considered in several papers about revision and inference (see e.g., [Rott, 2009]), but in a CWR each stratum must be a single formula and the JEPD condition must be satisfied. The aim is to compactly represent a total preorder over the interpretations.

We are interested in identifying the computational complexity of the following decision problem:

**Definition 6** ( $\models_{\star}$ ). *Let  $\star$  be a change operator.  $\models_{\star}$  is the decision problem defined as follows:*

- **Input:** A CWR  $\Phi$ , formulae  $\mu, \alpha$ .
- **Question:** Does  $B(\Phi \star \mu) \models \alpha$  hold?

We assume that the reader is familiar with the complexity classes NP and coNP. We briefly recall below some complexity classes of interest in this paper, all at the first level of the polynomial hierarchy. The class  $\Theta_2^P = \text{P}^{\text{NP}}[\mathcal{O}(\log n)]$  [Wagner, 1987; Eiter and Gottlob, 1997] is the class of languages that can be recognized in polynomial time by a deterministic Turing machine using a number of calls to an NP oracle bounded by a logarithmic function of the size of the input. The Boolean hierarchy **BH** is the Boolean closure of NP under Boolean operations (see [Cai *et al.*, 1988] for the formal definition of **BH**). Roughly speaking, a problem belongs to the class  $\text{NP}(k)$  or  $\text{coNP}(k)$  only if solving it requires at most  $k$  calls to an NP oracle.

### 3.1 Iterated Revision

Prop. 3 below tells us that for any DP-AGM operator  $\circ$ , one can precisely characterize the computational complexity of deciding whether a formula is entailed by a revised CWR: it “only” requires a logarithmic number of calls to an NP oracle. This result highly relies on the following lemma:

**Lemma 1.** *Let  $\circ$  be a DP-AGM operator. Then for any CWR  $\Phi = (\varphi_1, \dots, \varphi_n)$  and any formula  $\mu$ ,  $B(\Phi \circ \mu) \equiv \varphi_{i_*} \wedge \mu$ , where  $i_* = \min_{i \in \{1, \dots, n\}} \{i \mid \varphi_i \wedge \mu \not\models \perp\}$ .*

**Proposition 3.** *Let  $\circ$  be a DP-AGM operator. Then  $\models_{\circ}$  is  $\Theta_2^P$ -complete.*

These results echo results of the same nature about (one step) base revision [Nebel, 1998], and showing that the inference problem for revision is NP-hard and coNP-hard when the operator under consideration satisfies the revision postulates. They cohere with other results pointed out in [Liberatore, 1997], where the complexity of inference for a given sequence of change formulae is identified (the problem is  $\text{P}^{\text{NP}}$ -complete for many iterated revision operators).

### 3.2 Soft Improvement

Let us start with a useful lemma:

**Lemma 2.** *Let  $\bullet$  be a soft improvement operator,  $\Phi = (\varphi_1, \dots, \varphi_n)$  be a CWR and  $\mu$  be any formula. Then:*

- (a) *if  $\varphi_1 \wedge \mu \not\models \perp$  then  $B(\Phi \bullet \mu) \equiv \varphi_1 \wedge \mu$ ; and*

- (b) *if  $(\varphi_1 \vee \varphi_2) \wedge \mu \models \perp$  then  $B(\Phi \bullet \mu) \equiv \varphi_1$ .*

The above result provides a set of sufficient conditions to characterize the beliefs of a revised CWR: given a CWR  $\Phi = (\varphi_1, \dots, \varphi_n)$  and a formula  $\mu$ ,  $B(\Phi \bullet \mu)$  can be characterized in the case when  $\mu$  is consistent with  $\varphi_1$ , or when  $\mu$  is inconsistent with  $\varphi_1 \vee \varphi_2$ . As a consequence, one can derive the following hardness result for soft improvement operators:

**Proposition 4.** *Let  $\bullet$  be a soft improvement operator. Then  $\models_{\bullet}$  is coNP(3)-hard, even when  $\Phi$  has at most three strata.*

This means that for any soft improvement operator  $\bullet$ , at least three calls to an NP oracle are necessary to decide  $\models_{\bullet}$ . Interestingly, this is also an upper bound for the case of one-improvement. This result uses the following lemma which characterizes  $B(\Phi \bullet_O \mu)$  in the case not covered in Lemma 2:

**Lemma 3.** *Let  $\Phi = (\varphi_1, \dots, \varphi_n)$  be an ES,  $\mu$  be any formula, and assume that  $\varphi_1 \wedge \mu \models \perp$  and  $\varphi_2 \wedge \mu \not\models \perp$ . Then  $B(\Phi \bullet_O \mu) \equiv \varphi_1 \vee (\varphi_2 \wedge \mu)$ .*

**Proposition 5.** *Let  $\Phi = (\varphi_1, \dots, \varphi_n)$  and  $\alpha, \mu$  be two formulae. Then  $\models_{\bullet_O}$  can be decided by the following procedure:*

1. *if  $\varphi_1 \wedge \mu \wedge \neg \alpha \not\models \perp$  then return False;*
2. *if  $\varphi_1 \wedge \mu \not\models \perp$  then return True;*
3. *if  $(\varphi_1 \vee (\varphi_2 \wedge \mu)) \wedge \neg \alpha \not\models \perp$  then return False;*
4. *Return True;*

So as a consequence of Prop. 4 and 5, we get that:

**Proposition 6.**  *$\models_{\bullet_O}$  is coNP(3)-complete.*

For half-improvement and best-improvement, five calls to an NP oracle are necessary and sufficient to provide an answer to the inference problem: Prop. 7 and 8 below show membership to coNP(5), while Prop. 9 shows coNP(5)-hardness. The proofs take advantage of the following useful lemma, which shows how  $B(\Phi \bullet_H \mu)$  and  $B(\Phi \bullet_B \mu)$  are characterized in the remaining case not covered in Lemma 2:

**Lemma 4.** *Let  $\Phi = (\varphi_1, \dots, \varphi_n)$  be an ES,  $\mu$  be any formula, and assume that  $\varphi_1 \wedge \mu \models \perp$  and  $\varphi_2 \wedge \mu \not\models \perp$ . Then*

$$B(\Phi \bullet_H \mu) \equiv \begin{cases} \varphi_1 \vee \varphi_2 & \text{if } \varphi_2 \models \mu, \\ \varphi_1 & \text{otherwise.} \end{cases}$$

$$B(\Phi \bullet_B \mu) \equiv \begin{cases} \varphi_1 \vee \varphi_2 & \text{if } \mu \text{ is separated in } \preceq_{\Phi}, \\ \varphi_1 & \text{otherwise.} \end{cases}$$

**Proposition 7.** *Let  $\Phi = (\varphi_1, \dots, \varphi_n)$  and  $\alpha, \mu$  be two formulae. Then  $\models_{\bullet_H}$  can be decided by the following procedure:*

1. *if  $\varphi_1 \wedge \mu \wedge \neg \alpha \not\models \perp$  then return False;*
2. *if  $\varphi_1 \wedge \mu \not\models \perp$  then return True;*
3. *if  $\varphi_1 \wedge \neg \alpha \not\models \perp$  then return False;*
4. *if  $\varphi_2 \wedge \neg \mu \not\models \perp$  then return True;*
5. *if  $\varphi_2 \wedge \neg \alpha \not\models \perp$  then return False;*
6. *Return True;*

**Proposition 8.** *Let  $\Phi = (\varphi_1, \dots, \varphi_n)$  and  $\alpha, \mu$  be two formulae. Then  $\models_{\bullet_B}$  can be decided by the procedure described in Prop. 7, where line 4 is replaced by the following:*

4'. if  $\Gamma \not\models \perp$  then return True,

where  $\Gamma = \bigvee_{i=1}^n \varphi_i^{2i-1} \wedge \mu^{2i-1} \wedge \varphi_i^{2i} \wedge \neg \mu^{2i}$ , and each  $\varphi_i^j$  (resp.  $\mu^j$ ) is defined as the original formula  $\varphi_i$  (resp.  $\mu$ ) by renaming each propositional variable  $x$  into a fresh one  $x^j$ .

**Proposition 9.**  $\models_{\bullet_H}$  and  $\models_{\bullet_B}$  are coNP(5)-hard. Hardness holds in each case even when  $\Phi$  has at most three strata.

So from Prop. 7-9, we get that:

**Proposition 10.**  $\models_{\bullet_H}$  and  $\models_{\bullet_B}$  are coNP(5)-complete.

## 4 Implementing Iterated Belief Change

We describe now for each belief change operator  $\star$  among the seven considered in the previous section how to encode a changed CWR  $\Phi' = \Phi \star \mu$ , given any initial CWR  $\Phi = (\varphi_1, \dots, \varphi_n)$  and any change formula  $\mu$ . In every case,  $\Phi'$  consists of formulae that are Boolean combinations (obtained using the connectives  $\wedge, \vee, \neg$ ) of the input formulae  $\varphi_i$  and  $\mu$ . Noteworthy, our encodings could easily be adapted to other belief revision operators, e.g., those listed in [Rott, 2009].

### 4.1 Iterated Revision

For any DP-AGM operator  $\circ$ , the general idea in the construction of the CWR  $\Phi' = \Phi \circ \mu$  is described as follows. One searches for the smallest index  $i_*$  such that  $\varphi_{i_*} \wedge \mu$  is consistent, which defines the first stratum of  $\Phi \circ \mu$ . Then, depending on the operator  $\circ$  under consideration, the remaining strata  $\varphi_j$  ( $j \neq i_*$ ) of  $\Phi$  either (i) remain unchanged in  $\Phi \circ \mu$ , or (ii) are split into two consecutive strata  $\varphi_j \wedge \mu$  and  $\varphi_j \wedge \neg \mu$  in  $\Phi \circ \mu$ , or lastly (iii) are (disjunctively) merged with others and ranked according to an order specific to  $\circ$ .

In the descriptions, given a vector  $\Psi = (\psi_1, \dots, \psi_k)$  of formulae whose sets of models are JEPD (but where some formulae  $\psi_i$  may be inconsistent), we say that  $\Psi' = (\psi'_1, \dots, \psi'_l)$  is a  $\perp$ -filtering of  $\Psi$  when  $\Psi'$  is defined as the restriction of  $\Psi$  to its consistent formulae, listed in the same order (identifying those formulae requires consistency tests). Obviously enough, the resulting vector  $\Psi'$  is a CWR.

**Boutilier's natural revision.** We define the vector of formulae  $\Psi = (\psi_1, \dots, \psi_{n+1})$  as follows. We first search for  $i_*$  and define  $\psi_1 = \varphi_{i_*} \wedge \mu$ . Then we set  $\psi_{i+1} = \varphi_i$  for each  $i \in \{1, \dots, i_* - 1, i_* + 1, \dots, n\}$ , and  $\psi_{i_*+1} = \varphi_{i_*} \wedge \neg \mu$ . Then we define  $\Phi' = \Phi \circ_B \mu$  as the  $\perp$ -filtering of  $\Psi$ .

**Nayak's lexicographic revision.** We define the vector of formulae  $\Psi = (\psi_1, \dots, \psi_{2n})$  as  $\psi_i = \varphi_i \wedge \mu$  and  $\psi_{n+i} = \varphi_i \wedge \neg \mu$  for each  $i \in \{1, \dots, n\}$ . Then we define  $\Phi' = \Phi \circ_N \mu$  as the  $\perp$ -filtering of  $\Psi$ .

**Lehmann's ranked revision.** We define the vector of formulae  $\Psi = (\psi_1, \dots, \psi_{n-i_*+2})$  as  $\psi_1 = \varphi_{i_*} \wedge \mu$ ,  $\psi_2 = (\varphi_{i_*} \wedge \neg \mu) \vee \bigvee_{i=1}^{i_*-1} \varphi_i$ , and for each  $i \in \{3, \dots, n - i_* + 2\}$ ,  $\psi_i = \varphi_{i_*+i-2}$ . Then we define  $\Phi' = \Phi \circ_L \mu$  as the  $\perp$ -filtering of  $\Psi$ .

**Booth and Meyer's restrained revision.** We first define the vector of formulae  $\Psi = (\psi_1, \dots, \psi_{2n})$  as  $\psi_{2i-1} = \varphi_i \wedge \mu$  and  $\psi_{2i} = \varphi_i \wedge \neg \mu$  for each  $i \in \{1, \dots, n\}$ . Then, we associate with  $\Psi$  its  $\perp$ -filtering  $\Psi'$ . Lastly, we take advantage of the procedure described above for computing Boutilier's natural revision and define  $\Phi' = \Phi \circ_{BM} \mu$  as  $\Phi' = \Psi' \circ_B \mu$ .

## 4.2 Soft Improvement

In contrast to what happens in the iterated revision case, in order to perform an improvement of any formula into a CWR, we do not need to search for the corresponding  $\varphi_{i_*}$  in the CWR.

**One-improvement.** We define the vector of formulae  $\Psi = (\psi_1, \dots, \psi_{n+1})$  as  $\psi_1 = \varphi_1 \wedge \mu$ ,  $\psi_{n+1} = \varphi_n \wedge \neg \mu$ , and for each  $i \in \{2, \dots, n\}$ ,  $\psi_i = (\varphi_{i-1} \wedge \neg \mu) \vee (\varphi_i \wedge \mu)$ . Then we define  $\Phi' = \Phi \bullet_O \mu$  as the  $\perp$ -filtering of  $\Psi$ .

**Half-improvement.** We first define the vector of formulae  $\Psi = (\psi_1, \dots, \psi_{2n})$  as  $\psi_{2i-1} = \varphi_i \wedge \mu$  and  $\psi_{2i} = \varphi_i \wedge \neg \mu$ , for each  $i \in \{1, \dots, n\}$ . Then, we associate with  $\Psi$  the vector  $\Theta = (\theta_1, \dots, \theta_{2n})$  defined as  $\theta_1 = \psi_1$ ,  $\theta_{2n} = \varphi_{2n}$ , and for each  $i \in \{1, \dots, n-1\}$ ,

$$\begin{cases} [\theta_{2i} = \psi_{2i} \vee \psi_{2i+1} \text{ and } \theta_{2i+1} = \perp] & \text{if } \psi_{2i+2} \models \perp, \\ [\theta_{2i} = \psi_{2i} & \text{and } \theta_{2i+1} = \psi_{2i+1}] & \text{otherwise.} \end{cases}$$

Lastly, we define  $\Phi' = \Phi \bullet_H \mu$  as the  $\perp$ -filtering of  $\Theta$ .

**Best-improvement.** We first define the vector of formulae  $\Psi = (\psi_1, \dots, \psi_{2n})$  as  $\psi_{2i-1} = \varphi_i \wedge \mu$  and  $\psi_{2i} = \varphi_i \wedge \neg \mu$ , for each  $i \in \{1, \dots, n\}$ . Now, let  $\Psi'$  be the  $\perp$ -filtering of  $\Psi$ . Then  $\Phi' = \Phi \bullet_B \mu$  is defined as  $\Phi' = \Psi'$  if  $\Psi' \neq \Phi$ , otherwise  $\Phi' = \Psi' \bullet_O \mu$ .

## 5 Empirical Results

To evaluate from the practical viewpoint the belief change operators considered in the previous sections, we made some experiments. We generated belief change instances that consist of consistent, yet non-valid propositional formulae  $\varphi$  so that the CWR considered at start is of the form  $(\varphi, \neg \varphi)$ , and sequences of change formulae  $\mu$ , and we computed the corresponding changed CWR for each one of the seven operators. We measured the cumulated time required to compute the resulting CWR, and the corresponding number of strata. The propositional formulae  $\varphi$  and  $\mu$  considered in our experiments are 3-CNF formulae (i.e., conjunctions of clauses of size at most 3), built using the modularity-based random generator presented in [Giráldez-Cru and Levy, 2015]. This generator takes advantage of a modularity parameter, impacting the community structure of the produced instances. As advocated in this paper, the corresponding formulae can be considered as pseudo-industrial random instances given that industrial instances are not purely random but have some structure.

The SAT solver used in the experiments is MiniSAT [Eén and Sörensson, 2003]. It takes as inputs CNF formulae. Since the formulae occurring in the CWRs obtained after a sequence of change are not in CNF in the general case, achieving the consistency tests used to implement a change operation (cf. previous section) requires first to turn the corresponding formulae into CNF. This is done through the introduction of new variables, while preserving the set of logical consequences over the variables considered at start [Plaisted and Greenbaum, 1986]. All the experiments have been conducted on a cluster of Intel Xeon E5-2643 (3.30 GHz) quad core processors with 32 GiB RAM. A time-out of 900s and a memory-out of 7.6 GiB has been considered for each instance.

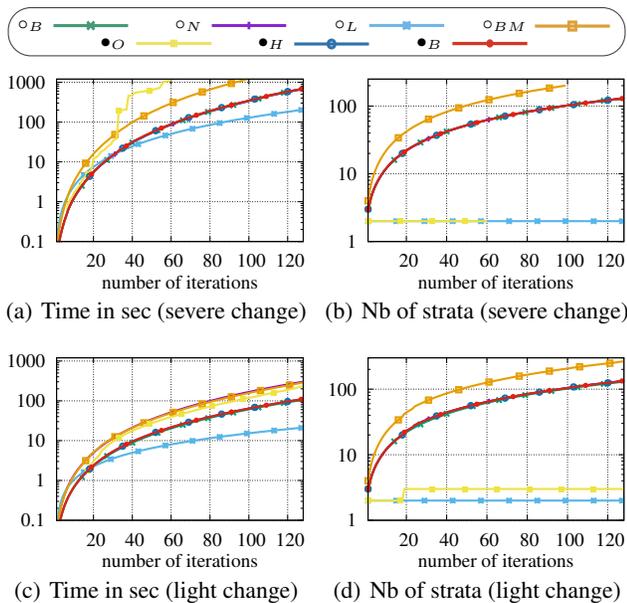


Figure 2: Empirical results.

Fig. 2(a) presents on its  $y$ -axis the median cumulated time (in seconds, over 100 instances) for computing the output CWRs associated with a sequence of change formulae, the length of it being given on the  $x$ -axis. For each instance, the 3-CNF formula  $\varphi$  considered at start is built over 4048 variables, has 12144 clauses, and is obtained using a modularity of 512, while every change formula  $\mu$  of the sequence is a 3-CNF formula generated from 814 variables (among the 4048 ones used in  $\varphi$ ), with 2442 clauses and for a modularity picked up at random between 3 and 50. Given the number of variables and clauses in it, every such  $\mu$  corresponds to a “severe” change scenario. It must also be noted that nothing ensures that the communities existing in any  $\mu$  are connected to the ones existing in  $\varphi$ . Fig. 2(b) gives on its  $y$ -axis the median number of strata in the resulting CWRs obtained for the same inputs, depending on the number of change formulae, given on the  $x$ -axis. Fig. 2(c) and 2(d) report results of the same nature as the ones given in Fig. 2(a) and 2(b), respectively, but considering “light” changes. Thus, every change formula  $\mu$  of the sequence is a 3-CNF formula generated from 45 variables (among the 4048 ones used in  $\varphi$ ), with 135 clauses. Logarithmic scales are used for the  $y$ -axis in each figure.

These empirical results show that both the time needed to compute the revised CWRs and the number of strata in them do not grow very quickly with the number of changes (unsurprisingly for Lehmann’s ranked revision, the number of strata never increases by definition). This can be explained by the fact that the logical strength of the formulae corresponding to the strata of a CWR increases with the number of strata. Indeed, though the complexity of achieving a change step depends on the number of consistency checks that are performed (hence, on the number of strata of the CWR), the consistency checks become at some point easier when the formulae of the CWR involved in them have few models, and are more likely to be inconsistent with the change formula (and as

such, filtered out). Experimentally, the nature of the change (severe vs. light) does not seem to have a strong impact on the number of strata in the resulting base (it has a more salient influence on the computation times). Interestingly, these experiments show that feasibility of iterated belief change for bases and sequences of change formulae of significant sizes.

## 6 Related Work

A closely related work is [Liberatore, 1997] where the complexity of iterated belief revision was investigated, given a sequence of change formulae. Our work differs from Liberatore’s one in a number of directions. First, among the seven operators we considered, five were not studied in [Liberatore, 1997] as they have been introduced a decade later. Second, some of our complexity results apply to any DP-AGM operator (Prop. 3) and any soft improvement operator (Prop. 4), while Liberatore [1997] focused on “concrete” revision operators. Lastly, the results from Liberatore [1997] are all theory-oriented, while we also performed experiments to show the extent to which iterated belief change is feasible in practice.

From the practical side, though there exist a couple of implementations of (one-step) belief revision operators (see mainly [Chou and Winslett, 1991; Dixon and Wobcke, 1993; Liberatore, 1999; Williams and Sims, 2000; Gorogiannis and Ryan, 2002; Delgrande *et al.*, 2007; Thimm, 2014; Konieczny *et al.*, 2017]), only few papers describe pieces of software for iterated belief revision. Indeed, the case of iterated belief revision is more tricky since it requires the representation of (complex) epistemic states. The system for iterated belief revision presented in [Zhuang *et al.*, 2007] is based on a notion of compiled epistemic entrenchment (roughly, each cluster of formulae in the ranking is turned into its prime implicate form). Such compiled epistemic entrenchments can be viewed as compact encodings of epistemic states based on formulae, which contrasts with the representations of epistemic states we considered (our CWRs are compact, formula-based encodings of epistemic states based on worlds). The experiments reported in [Zhuang *et al.*, 2007] consider a limited number of propositional variables (only 5).

## 7 Conclusion

In this paper, iterated belief change has been considered from a computational perspective. A number of complexity results for the inference problem has been provided for several (families of) iterated change operators, and empirical results have been reported as well. As future work, we plan to identify the computational complexity of the inference problem when a sequence of change formulae is considered, focusing on the change operators not considered in [Liberatore, 1997]. In addition, since iterated contraction operators [Booth and Chandler, 2016; Konieczny and Pino Pérez, 2017] are closely related to iterated revision operators, one can take advantage of our framework to encode and compute the resulting CWRs for these operators too. This is left for future work. Lastly, we will implement some other change operators (especially, those described in [Rott, 2009]) and make the code we developed for generating and querying CWRs available online.

## References

- [Alchourrón *et al.*, 1985] C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
- [Booth and Chandler, 2016] R. Booth and J. Chandler. Extending the harper identity to iterated belief change. In *IJCAI'16*, pages 987–993, 2016.
- [Booth and Meyer, 2006] R. Booth and T. A. Meyer. Admissible and restrained revision. *Journal of Artificial Intelligence Research*, 26:127–151, 2006.
- [Boutilier, 1996] C. Boutilier. Iterated revision and minimal change of conditional beliefs. *Journal of Philosophical Logic*, 25(3):263–305, 1996.
- [Cai *et al.*, 1988] J-Y. Cai, T. Gundermann, J. Hartmanis, L. A. Hemachandra, V. Sewelson, K. Wagner, and G. Wechsung. The Boolean hierarchy I: Structural properties. *SIAM Journal of Computing*, 17(6):1232–1252, 1988.
- [Chou and Winslett, 1991] S-C. T. Chou and M. Winslett. The implementation of a model-based belief revision system. *SIGART Bulletin*, 2(3):28–34, 1991.
- [Darwiche and Pearl, 1997] A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89(1-2):1–29, 1997.
- [Delgrande *et al.*, 2007] J. P. Delgrande, D. H. Liu, T. Schaub, and S. Thiele. COBA 2.0: A consistency-based belief change system. In *ECSQARU'07*, pages 78–90, 2007.
- [Dixon and Wobcke, 1993] S. Dixon and W. Wobcke. The implementation of a first-order logic AGM belief revision system. In *ICTAI'93*, pages 40–47, 1993.
- [Eén and Sörensson, 2003] N. Eén and N. Sörensson. An extensible SAT-solver. In *SAT'03*, pages 502–518, 2003.
- [Eiter and Gottlob, 1997] T. Eiter and G. Gottlob. The complexity class  $\theta_2^P$ : Recent results and applications in AI and modal logic. In *FCT'97*, pages 1–18, 1997.
- [Fermé and Hansson, 2011] E. Fermé and S. O. Hansson. AGM 25 years: Twenty-five years of research in belief change. *Journal of Philosophical Logic*, 40(2):295–331, 2011.
- [Gärdenfors, 1988] P. Gärdenfors. *Knowledge in flux*. MIT Press, 1988.
- [Giráldez-Cru and Levy, 2015] J. Giráldez-Cru and J. Levy. A modularity-based random SAT instances generator. In *IJCAI'15*, pages 1952–1958, 2015.
- [Gorogiannis and Ryan, 2002] N. Gorogiannis and M. Ryan. Implementation of belief change operators using BDDs. *Studia Logica*, 70(1):131–156, 2002.
- [Hansson, 1999] S. O. Hansson. *A Textbook of Belief Dynamics. Theory Change and Database Updating*. Kluwer, 1999.
- [Jin and Thielscher, 2007] Y. Jin and M. Thielscher. Iterated belief revision, revised. *Artificial Intelligence*, 171(1):1–18, 2007.
- [Konieczny and Pino Pérez, 2008] S. Konieczny and R. Pino Pérez. Improvement operators. In *KR'08*, pages 177–187, 2008.
- [Konieczny and Pino Pérez, 2017] S. Konieczny and R. Pino Pérez. On iterated contraction: Syntactic characterization, representation theorem and limitations of the levi identity. In *SUM'17*, pages 348–362, 2017.
- [Konieczny *et al.*, 2010] S. Konieczny, M. M. Grespan, and R. Pino Pérez. Taxonomy of improvement operators and the problem of minimal change. In *KR'10*, pages 161–170, 2010.
- [Konieczny *et al.*, 2017] S. Konieczny, J-M. Lagniez, and P. Marquis. Boosting distance-based revision using SAT encodings. In *LORI'17*, pages 480–496, 2017.
- [Lehmann, 1995] D. Lehmann. Belief revision, revised. In *IJCAI'95*, pages 1534–1540, 1995.
- [Liberatore, 1997] P. Liberatore. The complexity of iterated belief revision. In *ICDT'97*, pages 276–290, 1997.
- [Liberatore, 1999] P. Liberatore. BReLS: a system for revising, updating, and merging knowledge bases. In *NRAC'99*, 1999.
- [Nayak *et al.*, 2003] A. C. Nayak, M. Pagnucco, and P. Pappas. Dynamic belief revision operators. *Artificial Intelligence*, 146(2):193–228, 2003.
- [Nebel, 1998] B. Nebel. *How Hard is it to Revise a Belief Base?*, pages 77–145. Springer Netherlands, Dordrecht, 1998.
- [Plaisted and Greenbaum, 1986] D. A. Plaisted and S. Greenbaum. A structure-preserving clause form translation. *Journal of Symbolic Computation*, 2(3):293–304, 1986.
- [Rott, 2006] H. Rott. Revision by comparison as a unifying framework: Severe withdrawal, irrevocable revision and irrefutable revision. *Theoretical Computer Science*, 355(2):228–242, 2006.
- [Rott, 2009] H. Rott. Shifting priorities: Simple representations for twenty-seven iterated theory change operators. In D. Makinson, J. Malinowski, and H. Wansing, editors, *Towards Mathematical Philosophy*, pages 269–296. Springer, 2009.
- [Thimm, 2014] M. Thimm. Tweety: A comprehensive collection of java libraries for logical aspects of artificial intelligence and knowledge representation. In *KR'14*, 2014.
- [Wagner, 1987] K. W. Wagner. More complicated questions about maxima and minima, and some closures of NP. *Theoretical Computer Science*, 51:53–80, 1987.
- [Williams and Sims, 2000] M-A. Williams and A. Sims. SATEN: an object-oriented web-based revision and extraction engine. *CoRR*, cs.AI/0003059, 2000.
- [Zhuang *et al.*, 2007] Z. Q. Zhuang, M. Pagnucco, and T. Meyer. Implementing iterated belief change via prime implicates. In *AI'07*, pages 507–518, 2007.