

Translation-based Revision and Merging for Minimal Horn Reasoning

Gerhard Brewka¹ and Jean-Guy Mailly² and Stefan Woltran²

Abstract. In this paper we introduce a new approach for revising and merging consistent Horn formulae under minimal model semantics. Our approach is translation-based in the following sense: we generate a propositional encoding capturing both the syntax of the original Horn formulae (the clauses which appear or not in them) and their semantics (their minimal models). We can then use any classical revision or merging operator to perform belief change on the encoding. The resulting propositional theory is then translated back into a Horn formula. We identify some specific operators which guarantee a particular kind of minimal change. A unique feature of our approach is that it allows us to control whether minimality of change primarily relates to the syntax or to the minimal model semantics of the Horn formula. We give an axiomatic characterization of minimal change on the minimal model for this new setting, and we show that some specific translation-based revision and merging operators satisfy our postulates.

1 Introduction

Belief revision is a highly active area of research in the field of knowledge representation and reasoning. Whereas the initial focus was on revising propositional belief sets, as in the famous AGM theory [1], respectively propositional knowledge bases, as in the KM theory [20], later on also revision operators for various more expressive logics have been studied. Examples are description logics, e.g. [24], modal and multi-valued logics [16], and others. Interestingly, in recent years there has as well been a steadily growing interest in the opposite direction, that is, in fragments of propositional logic which are less expressive but interesting for some specific reason.

Belief change operators for fragments of propositional logic have been thoroughly studied in [9, 7, 8]. Horn formulae, *i.e.* conjunctions (or equivalently sets) of clauses which contain at most one positive literal [10, 19], play a special role in this context. Horn formulae are particularly interesting for computational reasons, as they allow for linear inference methods.

Revising Horn formulae is also the topic of this paper. However, contrary to the existing work cited above which considers classical reasoning based on *all* models of a Horn formula, we are interested here in reasoning under the *minimal model semantics*, or minimal reasoning, for short. This form of reasoning is based on the assumption that an atom should be considered *false* whenever it is not provably *true*. It is well-known that consistent Horn formulae have a unique minimal (and thus least) model. Since new information may

modify the least model in arbitrary ways, minimal reasoning is non-monotonic, and the revision operators we are looking for are actually operators for a nonmonotonic formalism.

Although the relationship between nonmonotonic reasoning and revision was already investigated by Gärdenfors [18], there is relatively little work on revising specific nonmonotonic formalisms. Notable exceptions are revision of default logic [3], revision of logic programs under answer set semantics [2, 15, 12, 11, 25, 26] and, rather recently, revision of argumentation frameworks [5, 4, 13].

The revision of logic programs is of special interest here. Since sets of Horn clauses under minimal model semantics are a special case of logic programs under stable semantics, one might say the problem is already solved. However, the work presented here is very different, as we will see. We will analyze these issues in detail in the discussion section.

The approach we are going to introduce in this paper is translation-based. The idea underlying translation-based revision has been pioneered in [17]. Basically, a revision operator for a formula F in a logic L is defined by representing F , possibly together with relevant meta-information, in a logic L^* for which a revision operator already exists. The encoding of F is then revised in L^* , and the result of this revision is translated back into logic L . Mailly [23] has shown how to apply this form of revision to a nonmonotonic formalism, namely to Dung-style argumentation frameworks [14].

In a nutshell, the goal of this paper is to demonstrate that translation-based revision can be also used for other, less simple nonmonotonic formalisms, in our case Horn formulae under minimal model semantics. In addition, we will also show how to use the translation-based approach for merging, that is, the process of integrating several Horn formulae into a single one.

More specifically, we encode the relation between the syntax of a Horn formula (its set of Horn clauses) and its semantics (its minimal model) in propositional logic. Revision or merging are then performed on this logical encoding with classical operators [20, 21], followed by a decoding step which gives the result of the Horn revision or merging. Our approach permits to revise a consistent Horn formula by an expressive piece of information: our revision formulae concern the set of clauses and the minimal model *at the same time*. Similarly, we use such an expressive logical language to express integrity constraints in the merging process. In contrast to other works, we thus do *not* restrict the revision formulae to the Horn fragment.

Let us schematically explain our approach to revision using Figure 1. Here, φ is a Horn formula, and $enc(\varphi)$ is the encoding of φ . We want to define a revision operator op for minimal model reasoning. We define our new revision operator op through three steps: encoding φ in propositional logic, resulting in $enc(\varphi)$, revising the encoding using an existing propositional operator op' , and finally de-

¹ Computer Science Institute, Leipzig University, Germany, brewka@informatik.uni-leipzig.de

² Institute of Information Systems, TU Wien, Austria, {jmailly,woltran}@dbai.tuwien.ac.at

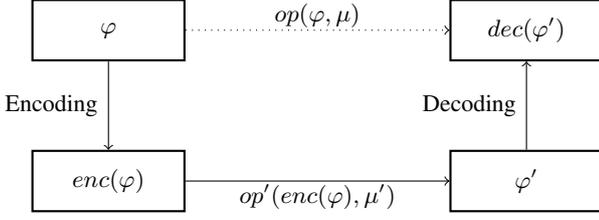


Figure 1. Schematic view of translation-based revision

coding the result of this revision back into a Horn formula. Note that the minimal Horn revision with μ requires revision with a substantially stronger formula μ' at the level of the encoding.

The schema for merging is similar to the one for revision. However, rather than a single Horn formula φ , a collection of such formulae (usually called a profile in the context of merging) is the starting point of the process, formulae μ and μ' can be seen as constraints for IC-merging [21] or are obsolete in simpler forms of merging.

The main contributions of this paper are the following. We introduce a propositional encoding of a consistent Horn formula, or equivalently set of clauses, φ such that the minimal model of φ is contained in the single model of the encoding. We use this encoding to define translation-based revision and merging of Horn formulae under minimal model semantics. For both operations,

- we define concrete translation-based operators which exhibit different minimal change properties;
- we adapt rationality postulates from the classical setting to express minimal change on the minimal model;
- we prove that some specific translation-based operators satisfy these postulates.

The remainder of the paper is organized as follows. Sect. 2 recalls the relevant background the paper builds upon. Sect. 3 provides the encoding of Horn formulae under minimal model semantics used throughout the paper. Sect. 4 and Sect. 5 introduce and analyze the translation-based approach to revision and merging, respectively. Sect. 6 shows how to limit the size of the encoding – which is exponential at the conceptual level – in practical settings. Sect. 7 discusses related work and concludes the paper with an outlook on future work. Some of the proofs are given in an appendix.

2 Background Notions

Let us first recall the basic notions of propositional logic and introduce some notations. We consider a set of Boolean variables V . We use \mathcal{L} to denote the set of all propositional formulae built on V with usual connectives (\neg, \vee, \wedge). Satisfiability of formulae is defined as usual, and $\text{mod}(\varphi)$ gives the set of models of a formula φ . Interpretations and models are represented by sets containing those variables which are assigned true. Each variable $x \in V$ is associated with a positive literal x and a negative literal $\neg x$. A clause is a disjunction of literals $l_1 \vee l_2 \vee \dots \vee l_n$. A Horn clause is a clause which contains at most one positive literal. $\text{hcl}(V)$ is the set of all Horn clauses built on V . A Horn formula φ is a conjunction of Horn clauses, or equivalently a set of Horn clauses, denoted $\text{hcl}(\varphi)$. $\mathcal{L}_H(V)$ denotes the set of all Horn formulae built on V , $\mathcal{L}_H^{\text{cons}}(V)$ the subset of consistent formulae in $\mathcal{L}_H(V)$. Given the set of models $\text{mod}(\varphi)$ of a formula $\varphi \in \mathcal{L}_H^{\text{cons}}(V)$, the minimal model of φ , denoted $\text{mod}_{\min}(\varphi)$,

is the unique \subseteq -minimal element of $\text{mod}(\varphi)$. $\text{mod}_{\min}(\varphi)$ models skepticism regarding positive atomic information, since each variable which is assigned *true* in this model is also assigned *true* in each other model of φ .

Given two sets S_1, S_2 , the Hamming distance between them is $d_H = |(S_1 \setminus S_2) \cup (S_2 \setminus S_1)|$. When S_1 and S_2 are propositional interpretations, the Hamming distance counts the number of Boolean variables which are assigned different values in these interpretations. Given a set of Boolean variables S , d_H^S is the Hamming distance between interpretations restricted to S .

Belief revision incorporates a new piece of information in an agent's beliefs. One of the most prominent characterizations of belief revision considers the beliefs and the new piece of information as formulae from propositional logic [20]. An axiomatic characterization is given by a set of postulates which express logical relations between formulae. We give here a reformulation of these postulates as set-theoretical relations between the sets of models of the formulae. A KM revision operator \circ is a mapping from $\mathcal{L} \times \mathcal{L}$ to \mathcal{L} which satisfies the postulates:

- (R1) $\text{mod}(\varphi \circ \mu) \subseteq \text{mod}(\mu)$.
- (R2) If $\text{mod}(\varphi) \cap \text{mod}(\mu) \neq \emptyset$, then $\text{mod}(\varphi \circ \mu) = \text{mod}(\varphi) \cap \text{mod}(\mu)$.
- (R3) If $\text{mod}(\mu) \neq \emptyset$, then $\text{mod}(\varphi \circ \mu) \neq \emptyset$.
- (R4) If $\text{mod}(\varphi_1) = \text{mod}(\varphi_2)$ and $\text{mod}(\mu_1) = \text{mod}(\mu_2)$, then $\text{mod}(\varphi_1 \circ \mu_1) = \text{mod}(\varphi_2 \circ \mu_2)$.
- (R5) $\text{mod}(\varphi \circ \mu_1) \cap \text{mod}(\mu_2) \subseteq \text{mod}(\varphi \circ (\mu_1 \wedge \mu_2))$.
- (R6) If $\text{mod}(\varphi \circ \mu_1) \cap \text{mod}(\mu_2) \neq \emptyset$, then $\text{mod}(\varphi \circ (\mu_1 \wedge \mu_2)) \subseteq \text{mod}(\varphi \circ \mu_1) \cap \text{mod}(\mu_2)$.

These postulates express logical constraints which must be satisfied by the models of the revised formula $\varphi \circ \mu$, depending on the models of the initial formula φ and the revision formula μ . A representation theorem associates this axiomatic characterization with a constructive one: a revision operator \circ satisfies the postulates iff it can be expressed as

$$\text{mod}(\varphi \circ \alpha) = \min(\text{mod}(\alpha), \leq_{\varphi}) \quad (1)$$

where \leq_{φ} is a total pre-order expressing the relative plausibility of interpretations; this pre-order has to satisfy some conditions. We only exhibit a specific family of pre-orders based on distances. Given a distance d , we overload the notation and define $d(\omega, \omega') = \min_{\omega'' \in \text{mod}(\varphi)} d(\omega, \omega'')$. The pre-order \leq_{φ}^d is then defined by $\omega_1 \leq_{\varphi}^d \omega_2$ iff $d(\omega_1, \varphi) \leq d(\omega_2, \varphi)$. Instantiating equation (1) with $\leq_{\varphi} = \leq_{\varphi}^d$ defines a revision operator which satisfies all rationality postulates.

Belief merging operations obtain the beliefs of a group from the beliefs of each member of the group. In particular, belief merging with integrity constraints (IC-merging) is a generalization of belief revision [21]. In this scenario, we want to merge a tuple of formulae $E = \langle \varphi_1, \dots, \varphi_n \rangle$ called a *profile*, where each formula represents an agent's belief. It is expected that the result of the merging satisfies an integrity constraint μ . The result of such an operation is denoted $\Delta_{\mu}(E)$. Revision and IC-merging are strongly connected; indeed when IC-merging is performed on a single formula, the result yields a revision operation, i.e. there is a KM revision operator \circ such that $\Delta_{\mu}(\langle \varphi \rangle) \equiv \varphi \circ \mu$. Rationality postulates (which generalize the postulates for revision from above) and a representation theorem have also been stated for IC-merging. Similarly to revision operators, an IC-merging operator can be defined thanks to a total pre-order which represents the relative plausibility of interpretations: $\text{mod}(\Delta_{\mu}(E)) = \min(\text{mod}(\mu), \leq_E)$. The pre-order should satisfy

additional properties. We focus here on the method to define a pre-order from a distance. We need first the notion of aggregation functions. An *aggregation function* is a function \otimes which associates a non-negative real number to every finite tuple of non-negative numbers, and which satisfies:

- if $y \leq z$ then $\otimes(x_1, \dots, y, \dots, x_n) \leq \otimes(x_1, \dots, z, \dots, x_n)$;
- $\otimes(x_1, \dots, x_n) = 0$ iff $x_1 = \dots = x_n = 0$;
- $\forall x \in \mathbb{R}^+, \otimes(x) = x$.

For any distance d and any profile $E = \langle \varphi_1, \dots, \varphi_n \rangle$, we overload the notation of distances: $d(\omega, E) = \otimes(d(\omega_1, \varphi_1), \dots, d(\omega_n, \varphi_n))$. Now, given a distance d and an aggregation function \otimes , the pre-order $\leq_E^{d, \otimes}$ is defined by $\omega_1 \leq_E^{d, \otimes} \omega_2$ iff $d(\omega_1, E) \leq d(\omega_2, E)$.

3 Encoding Horn Formulae

The principle of our encoding is to define a propositional formula Ξ which establishes the links between the syntax of a consistent Horn formula, represented by the set of clauses which appear in it, and the semantics of the formula, as given by its minimal model. Then we need a way to encode the syntax of a Horn formula φ , such that the conjunction of this encoding with Ξ permits to deduce the minimal model of φ .

Let us introduce the propositional encoding of the syntax of a formula.

Definition 1. Let $\varphi \in \mathcal{L}_H^{cons}(V)$. For each Horn clause c built on V , a fresh variable cl_c is introduced with the intended meaning that the clause c appears in φ . Then the syntax of φ is encoded by

$$CL(\varphi) = \left(\bigwedge_{c \in \text{hcl}(\varphi)} cl_c \right) \wedge \left(\bigwedge_{c \in \text{hcl}(V) \setminus \text{hcl}(\varphi)} \neg cl_c \right)$$

Let us notice that for each variable $x \in V$, cl_x is true iff the unit clause x appears in the formula.

Example 1. We describe the encoding of the formula $\varphi = (x_1 \vee \neg x_2 \vee \neg x_3) \wedge (\neg x_1 \vee x_2 \vee \neg x_3) \wedge x_1$. Here, $\text{hcl}(\varphi) = \{x_1 \vee \neg x_2 \vee \neg x_3, \neg x_1 \vee x_2 \vee \neg x_3, x_1\}$.

$$CL(\varphi) = cl_{x_1 \vee \neg x_2 \vee \neg x_3} \wedge cl_{\neg x_1 \vee x_2 \vee \neg x_3} \wedge cl_{x_1} \\ \wedge \left(\bigwedge_{c \in \text{hcl}(V) \setminus \text{hcl}(\varphi)} \neg cl_c \right)$$

Now we need a formula which expresses the link between the syntax of a consistent Horn formula and its minimal model. Note that below we do ignore purely negative clauses since they do not influence the computation of the minimal model for consistent formulae.

Definition 2. We consider Horn formulae in $\mathcal{L}_H^{cons}(V)$. Let $|V| = n$. The encoding of Horn minimal model semantics is given by the formula $H(V)$ defined as follows:

$$\begin{aligned} H^{(x,0)} &= x^0 \Leftrightarrow cl_x \\ H^{(x,i)} &= x^i \Leftrightarrow (x^{i-1} \vee ded^{(x,i)}), \text{ for } 1 \leq i < n, \text{ where} \\ ded^{(x,i)} &= \bigvee_{c = \neg y_1 \vee \dots \vee \neg y_k \vee x} (cl_c \wedge y_1^{i-1} \wedge \dots \wedge y_k^{i-1}) \\ H^{(x,n)} &= x^n \Leftrightarrow x \\ H(V) &= \bigwedge_{x \in V, 0 \leq i \leq n} H^{(x,i)} \end{aligned}$$

This formula mimics the well-known linear-time marking algorithm for computing the minimal model of a Horn formula. The variables $\{x^i \mid x \in V, 0 \leq i \leq n\}$ are used to represent the state of the variable x at the i^{th} step of the algorithm. The algorithm is guaranteed to terminate after at most $n = |V|$ steps. For this reason it suffices to consider subformulae $H^{(x,j)}$ for $j \leq n$.

They are initialized by $H^{(x,0)}$, which states that the variables x is true at the beginning of the algorithm iff it appears as a unit clause in φ . Then, from $H^{(x,i)}$ we obtain that the variable x is true at the i^{th} step iff either it was true at the previous step, or it can be deduced from the Horn clauses which appear in φ and the variables which were true at the $i-1^{th}$ step. This part is represented by $ded^{(x,i)}$.

Finally, $H^{(x,n)}$ states that the variables which are true at the n^{th} step are those which are true in the minimal model of φ .

Example 2. Let $\varphi = (x_1 \vee \neg x_2 \vee \neg x_3) \wedge (x_2 \vee \neg x_4) \wedge (x_3 \vee \neg x_4) \wedge x_4$, and consider an arbitrary model M of $CL(\varphi) \wedge H(V)$.

- $H^{(x,0)}$: for each $x \neq x_4$ we have $M(x^0) = 0$ since there is no unit clause x in φ (so cl_x is false);
- $H^{(x_4,0)}$: $M(x_4^0) = 1$ since there is a unit clause x_4 in φ (so cl_{x_4} is true);
- $H^{(x_2,1)}, H^{(x_3,1)}$: at step 1, two rules from $ded^{(x,1)}$ can be applied (since x_4 allows to trigger them). So, for each $x \in \{x_2, x_3\}$, $M(x^1) = 1$.
- $H^{(x_1,1)}, H^{(x_4,1)}$: x_4^1 receives the value 1 because $M(x_4^0) = 1$; x_1^1 receives the value 0 because no rule concerning it can be triggered.
- $H^{(x,2)}$: at step 2, since x_2 and x_3 are true from the previous step, a rule is triggered and $M(x_1^2) = 1$. For other variables $x \in \{x_2, x_3, x_4\}$, $M(x^2) = 1$ because $M(x^1) = 1$.
- The same scheme is repeated for each i : $M(x^i) = 1$ because $M(x^{i-1}) = 1$. Finally, from $H^{(x,n)}$, each variable x_1, \dots, x_4 receives the value 1.

$\text{mod}_{\min}(\varphi) = \{x_1, x_2, x_3, x_4\}$ can thus be deduced from $CL(\varphi) \wedge H(V)$.

We can show the following result:

Proposition 1. Let $\varphi \in \mathcal{L}_H^{cons}(V)$. The propositional formula $CL(\varphi) \wedge H(V)$ has a unique model M such that $M \cap V = \text{mod}_{\min}(\varphi)$.

The reader will have noticed that the size of our encoding is actually exponential in the number of propositional variables in V . We are fully aware that this is far from tolerable from a practical point of view. For the time being we will stick to this exponential encoding, as this makes it easier to introduce our approach at the conceptual level. However, we will discuss in Sect. 6 how to deal with this issue in practical settings.

4 Revising Horn Formulae

We focus here on belief revision in a situation where the relevant information is carried by the minimal model of formulae. We call this operation mod_{\min} -revision.

4.1 Translation-based Revision

Our approach benefits from the encoding presented in the previous section. Indeed, the formula $CL(\varphi) \wedge H(V)$ expresses information about the minimal model of φ , which is given by the value of the variables $V = \{x_1, \dots, x_p\}$. It is thus possible to revise $CL(\varphi) \wedge H(V)$ by a propositional formula μ built on V which expresses the new piece of information to incorporate in the minimal model of the agent's beliefs. On the other hand, $CL(\varphi) \wedge H(V)$ also expresses information about the structure of the formula φ , since $CL(\varphi)$ is built using the cl_x variables. So we can include

structural information as well and use revision formulae μ in \mathcal{L}' , the propositional language built on $V' = V \cup \{\text{cl}_c \mid c \in \text{hcl}(V)\}$.

One important issue needs to be addressed: performing the revision $(CL(\varphi) \wedge H(V)) \circ \mu$, with \circ any KM revision operator, does not guarantee to give a result which is compatible with the Horn minimal model semantics. As a simple example, if we revise a formula φ by $\mu = x_1 \wedge \neg x_2 \wedge \text{cl}_{\neg x_1 \vee x_2}$, then only x_1 should belong to the minimal model of the result, and the clause $\neg x_1 \vee x_2$ should appear in this result. This is obviously incompatible. For this reason, we define revision operators as follows.

Definition 3. Let \circ be an arbitrary KM revision operator, $\varphi \in \mathcal{L}_H^{\text{cons}}(V)$ and $\mu \in \mathcal{L}'$. The translation-based revision operator based on \circ , denoted \star_\circ , is a mapping from $\mathcal{L}_H^{\text{cons}}(V) \times \mathcal{L}'$ to $2^{\mathcal{L}_H^{\text{cons}}(V)}$, defined as:

$$\varphi \star_\circ \mu = \text{dec}((CL(\varphi) \wedge H(V)) \circ (\mu \wedge H(V)))$$

with dec the decoding of the clause variables cl_x , defined as follows. Let ω be a propositional interpretation built on $V' = V \cup \{\text{cl}_c \mid c \in \text{hcl}(V)\}$. Let Ω be a set of such interpretations. Let Φ be a propositional formula built on V' .

- $\text{dec}(\omega)$ is the Horn formula $\{c \in \text{hcl}(V) \mid \omega(\text{cl}_c) = 1\}$;
- $\text{dec}(\Omega) = \{\text{dec}(\omega) \mid \omega \in \Omega\}$;
- $\text{dec}(\Phi) = \text{dec}(\text{mod}(\Phi))$.

Since \circ satisfies the KM rationality postulates, the result of the revision is obviously inconsistent whenever μ is not consistent with $H(V)$, as intended.

Now we define different revision operators \star . In particular, we show that – depending on the underlying operator \circ – we can actually choose between minimizing change on the semantic level of a formula (its minimal model) or on the syntactic level (the set of clauses of a formula).

Definition 4. Given $W_1 > 0, W_2 > 0$, we define the (W_1, W_2) -weighted distance between interpretations $d_{(W_1, W_2)}$ as follows:

$$d_{(W_1, W_2)}(\omega_1, \omega_2) = W_1 \times d_H^V(\omega_1, \omega_2) + W_2 \times d_H^{\text{Syn}}(\omega_1, \omega_2)$$

with ω_1, ω_2 interpretations on the set of variables $V \cup \{\text{cl}_c \mid c \in \text{hcl}(V)\}$, and $\text{Syn} = \{\text{cl}_c \mid c \in \text{hcl}(V)\}$ the set of variables related to syntax of formulae.

Here we consider only the variables which correspond to the syntax (cl_c variables) and to the minimal model semantics (x variables). But our encoding also uses additional variables x^i . For this reason, we cannot directly define a distance-based revision operator from $d_{(W_1, W_2)}$ as explained in Section 2, since $d_{(W_1, W_2)}$ is not strictly speaking a distance between interpretations of our encoding. However, we prove that we can use $d_{(W_1, W_2)}$ to define a KM revision operator.

Proposition 2. The pre-order between interpretations corresponding to $d_{(W_1, W_2)}$ satisfies the properties of faithful assignments, and yields a KM revision operator denoted $\circ_{(W_1, W_2)}$.

This family of weighted distances is a generalization of the Hamming distance ($d_{(1,1)}$ yields Hamming distance). It is possible to assign particular weights to obtain some properties about minimal change; if the value of W_1 is high enough, then it is more expensive in the revision process to change the value of a x variable than to

change the values of all cl_c variables. Then the revision operator will ensure the minimal change on the x variables (which represent the minimal model of the Horn formula); minimal change of the syntax will be applied as a secondary criterion.

Definition 5. Let $W_1 = |\text{hcl}(V)| + 1$, and $W_2 = 1$. The semantic minimal change revision operator is the translation-based revision operator \star_{sem} based on the KM revision operator $\circ_{(W_1, W_2)}$.

Example 3. Given $V = \{x_1, x_2, x_3, x_4, x_5\}$, we revise $\varphi = (x_1 \vee \neg x_2 \vee \neg x_3) \wedge (\neg x_1 \vee x_2 \vee \neg x_3) \wedge x_1$ from Example 1 by $\mu = x_4 \vee x_5$. We can see that $\text{mod}_{\min}(\varphi) = \{x_1\}$. Minimal change of the minimal model, in this situation, leads to $\text{mod}_{\min}(\varphi \star \mu) = \{\{x_1, x_4\}, \{x_1, x_5\}\}$. For each of these possible minimal models there is a possible Horn formula which corresponds to it and which is minimal with respect to the secondary criterion: $\varphi_1 = \varphi \wedge x_4$, $\text{mod}_{\min}(\varphi_1) = \{x_1, x_4\}$; and $\varphi_2 = \varphi \wedge x_5$, $\text{mod}_{\min}(\varphi_2) = \{x_1, x_5\}$. For each of these formulae, there is only one new variable in the minimal model, and one new clause in the formula. So the result of the revision is $\varphi \star_{\text{sem}} \mu = \{\varphi_1, \varphi_2\}$.

In Example 3, all the solutions are optimal with respect to minimal change of minimal model and minimal change of the syntax, but it is not the case in general. If it is more expensive to change the value of a single cl_c variable than to change the values of all x variables, then minimal change of the syntax is the main minimality criterion, and minimal change of the minimal model is applied as a secondary criterion.

Definition 6. Let $W_1 = 1$, and $W_2 = |V| + 1$. The syntactic minimal change revision operator is the translation-based revision operator \star_{syn} based on the KM revision operator $\circ_{(W_1, W_2)}$.

Even when the formula μ concerns only one kind of information (the minimal model, or the syntax), both operators lead to a different result in general.

Example 4. We exemplify the difference between both kinds of minimal change. We consider the set of Boolean variables $V = \{x_1, x_2, x_3, x_4\}$ and the formula $\varphi = (x_1 \vee \neg x_2 \vee \neg x_3) \wedge (x_2 \vee \neg x_4) \wedge (x_3 \vee \neg x_4) \wedge x_4$ from Example 2. Its minimal model is $\text{mod}_{\min}(\varphi) = \{x_1, x_2, x_3, x_4\}$. We want to revise it by the formula $\mu = \neg x_2$, which means that x_2 should not belong to the minimal model.

The encoding of φ is

$$CL(\varphi) = \text{cl}_{x_1 \vee \neg x_2 \vee \neg x_3} \wedge \text{cl}_{x_2 \vee \neg x_4} \wedge \text{cl}_{x_3 \vee \neg x_4} \wedge \text{cl}_{x_4} \\ \wedge \bigwedge_{c \in \text{hcl}(V) \setminus \text{hcl}(\varphi)} \neg \text{cl}_c$$

Horn minimal model semantics is encoded by $H(V) = \bigwedge_{x \in V, 0 \leq i \leq 4} H^{(x, i)}$.

We perform first a revision with the semantic minimal change operator \star_{sem} . Obviously, giving priority to minimal change on the minimal model leads to removing x_2 from it, which means that $\text{mod}_{\min}(\varphi \star_{\text{sem}} \mu) = \{x_1, x_3, x_4\}$. There are four possible Horn formulae, corresponding to this minimal model, which are minimal w.r.t. syntax change: $\varphi \star_{\text{sem}} \mu = \{(x_1 \vee \neg x_2 \vee \neg x_3) \wedge (x_3 \vee \neg x_4) \wedge x_4 \wedge X \mid X \in \{x_1, x_1 \vee \neg x_4, x_1 \vee \neg x_3, x_1 \vee \neg x_3 \vee \neg x_4\}\}$. In this case, there is a single change in the minimal model, and two changes in the syntax of the formula (one clause is removed, one clause is added). Alternatively, if we first consider minimal change of the syntax, then the result is the set of formulae $\varphi \star_{\text{syn}} \mu = \{(x_1 \vee \neg x_2 \vee \neg x_3) \wedge (x_3 \vee \neg x_4) \wedge x_4\}$, and $\text{mod}_{\min}(\varphi \star_{\text{syn}} \mu) = \{x_3, x_4\}$. Here, there is a single change in the syntax (the removal of one clause), and there are two changes in the minimal model (x_1 and x_2 are removed).

4.2 Axiomatization of mod_{\min} -Revision

If the minimal model of a formula is the most important information for an agent, then she can revise her beliefs by a formula μ which expresses what the new minimal model of her beliefs should be. The result of the revision is then a set of Horn formulae³, the minimal model of each of them being a model of μ . For a set of formulae Φ , we use $\text{mod}_{\min}(\Phi)$ as a notation for $\{\text{mod}_{\min}(\varphi) \mid \varphi \in \Phi\}$. This kind of revision should satisfy the following postulates:

- (★1) $\text{mod}_{\min}(\varphi \star \mu) \subseteq \text{mod}(\mu)$.
- (★2) If $\{\text{mod}_{\min}(\varphi)\} \cap \text{mod}(\mu) \neq \emptyset$, then $\text{mod}_{\min}(\varphi \star \mu) = \{\text{mod}_{\min}(\varphi)\} \cap \text{mod}(\mu)$.
- (★3) If $\text{mod}(\mu) \neq \emptyset$, then $\text{mod}_{\min}(\varphi \star \mu) \neq \emptyset$.
- (★4) If $\text{mod}_{\min}(\varphi_1) = \text{mod}_{\min}(\varphi_2)$ and $\text{mod}(\mu_1) = \text{mod}(\mu_2)$, then $\text{mod}_{\min}(\varphi_1 \star \mu_1) = \text{mod}_{\min}(\varphi_2 \star \mu_2)$.
- (★5) $\text{mod}_{\min}(\varphi \star \mu_1) \cap \text{mod}(\mu_2) \subseteq \text{mod}_{\min}(\varphi \star (\mu_1 \wedge \mu_2))$.
- (★6) If $\text{mod}_{\min}(\varphi \star \mu_1) \cap \text{mod}(\mu_2) \neq \emptyset$, then $\text{mod}_{\min}(\varphi \star (\mu_1 \wedge \mu_2)) \subseteq \text{mod}_{\min}(\varphi \star \mu_1) \cap \text{mod}(\mu_2)$.

The notion of faithful assignments in classical revision aims at sorting interpretations to ensure that the models of the agent's beliefs (which are the main information of the agent) are the minimal elements of a pre-order. In our case, we can relax this assumption, as soon as the minimal model of the formula is the "preferred" interpretation.

Definition 7. A min-faithful assignment is a mapping from a Horn formula φ to a total pre-order between interpretations \leq_{φ} such that:

1. $\forall \omega \neq \text{mod}_{\min}(\varphi), \text{mod}_{\min}(\varphi) <_{\varphi} \omega$;
2. if $\text{mod}_{\min}(\varphi_1) = \text{mod}_{\min}(\varphi_2)$, then $\leq_{\varphi_1} = \leq_{\varphi_2}$.

Theorem 3. Given a min-faithful assignment which maps each Horn formula φ to a total pre-order \leq_{φ} , if the revision operator \star satisfies $\text{mod}_{\min}(\varphi \star \mu) = \min(\text{mod}(\mu), \leq_{\varphi})$ then \star satisfies the postulates (★1)–(★6).

We notice that contrary to the classical theorem from Katsuno and Mendelzon, this result does not lead directly to the result for revision. Indeed, while the classical faithful assignment theorem characterizes the set of models of a formula (which is then unique, modulo logical equivalence), here we only know that the result of the revision should be a set of Horn formulae such that the minimal model of each of them belongs to $\min(\text{mod}(\mu), \leq_{\varphi})$. In general, for each minimal model $\omega \in \min(\text{mod}(\mu), \leq_{\varphi})$, there is not a single Horn formula which corresponds to ω . So, for a given min-faithful assignment, there may be several options to define the result of the revision.

We now focus on a particular family of revision operators, based on distances between interpretations.

Definition 8. Let d be a distance between interpretations. For each Horn formula φ , we define the total pre-order between interpretations \leq_{φ} by $\omega_1 \leq_{\varphi}^d \omega_2$ iff $d(\omega_1, \text{mod}_{\min}(\varphi)) \leq d(\omega_2, \text{mod}_{\min}(\varphi))$. The distance-based revision operator \star_d is defined by

$$\text{mod}_{\min}(\varphi \star_d \mu) = \min(\text{mod}(\mu), \leq_{\varphi}^d)$$

Proposition 4. Every distance-based revision operator satisfies the postulates (★1)–(★6).

³ If a specific application requires us to have a single Horn formula, a tie-break rule can be used on the result of the revision.

To illustrate distance-based revision, we look at a specific operator from this family.

Example 5. Let \star_H the distance based revision operator defined from the Hamming distance. We consider the Horn formula φ from Example 3 built on the set of variables $V = \{x_1, x_2, x_3, x_4, x_5\}$. Recall that $\text{mod}_{\min}(\varphi) = \{x_1\}$.

Now we revise φ by $\mu = x_4 \vee x_5$, which expresses that the minimal model of the agent's beliefs should contain at least one of the variables x_4, x_5 . The models of μ which are minimal w.r.t. the pre-order associated with φ and the Hamming distance are $M = \{\{x_1, x_4\}, \{x_1, x_5\}\}$. So the result of the revision should be a set of Horn formulae such that their minimal model belongs to M .

A possible solution to obtain the result of the revision is to consider the set of Horn formulae $R = \{\psi \mid \text{mod}_{\min}(\psi) \in \min(\text{mod}(\mu), \leq_{\varphi})\}$. This set represents all the possible alternatives for the agent's revised beliefs. Let us notice that this result can be refined thanks to the translation-based operators defined in Section 4.1.

Proposition 5. The semantic minimal change revision operator \star_{sem} of Definition 5 satisfies the postulates (★1)–(★6), provided the revision formula μ is built on V .

This translation-based revision operator satisfies the postulates, and is a concrete method to obtain the result of the revision (the set of revised Horn formulae), while the distance-based revision given in Definition 8 only gives the revised minimal models, but not the actual formulae corresponding to these models.

5 Merging Horn Formulae

Now we turn our attention to merging operators. Similarly to what we have proposed for revision, we propose operators such that not all models of formulae are considered in the merging process, but only their minimal model. We call this operation mod_{\min} -merging. To avoid confusion with classical IC-merging operators, we use Θ to denote mod_{\min} -merging operators. We focus on Horn profiles, meaning $\langle \varphi_1, \dots, \varphi_n \rangle$ such that $\varphi_1, \dots, \varphi_n \in \mathcal{L}_H^{cons}(V)$.

5.1 Translation-based Merging

We can also benefit from the logical encoding of syntax and mod_{\min} -semantics of Horn formulae to define merging operators. The idea is the same as for translation-based revision: we propose to translate all Horn formulae in the profile into propositional formulae dealing with the syntax and the minimal model, and to merge them with a classical IC-merging operator. The formula μ gives a constraint on the set of clauses and on the minimal model.

We first define the encoding of a Horn profile.

Definition 9. Given a Horn profile $E = \langle \varphi_1, \dots, \varphi_m \rangle$, the encoding of E is $H(E) = \langle CL(\varphi_1) \wedge H(V), \dots, CL(\varphi_m) \wedge H(V) \rangle$.

Definition 10. Given Δ any IC-merging operator, the translation-based merging operator based on Δ is Θ^{Δ} , a mapping from a Horn profile E and a formula $\mu \in \mathcal{L}^1$ to a set of consistent Horn formulae, such that $\Theta_{\mu}^{\Delta}(E) = \text{dec}(\Delta_{(\mu \wedge H(V))}(H(E)))$ with dec as given in Definition 3.

We use the well-known distance-based IC-merging operators to define mod_{\min} -merging operators. To exhibit specific operators, we use the sum as aggregation function.

Definition 11. Given the (W_1, W_2) -weighted distance between interpretations $d_{(W_1, W_2)}$, the (W_1, W_2) -weighted IC-merging operator $\Delta^{(W_1, W_2), \Sigma}$ is the distance based IC-merging operator defined from the distance $d_{(W_1, W_2)}$ and the sum aggregation function.

Now we define translation-based merging operators which are generalizations of the revision operators \star_{sem} and \star_{syn} .

Definition 12. Let $W_1 = |\text{hcl}(V)| + 1$, and $W_2 = 1$. The semantic minimal change merging operator is the translation-based merging operator Θ^{sem} based on the IC-merging operator $\Delta^{(W_1, W_2), \Sigma}$.

Definition 13. Let $W_1 = 1$, and $W_2 = |V| + 1$. The syntactic minimal change merging operator is the translation-based merging operator Θ^{syn} based on the IC-merging operator $\Delta^{(W_1, W_2), \Sigma}$.

Example 6. We consider $V = \{x_1, x_2, x_3, x_4\}$, and $E = \langle \varphi_1, \varphi_2 \rangle$, with $\varphi_1 = (x_1 \vee \neg x_2 \vee \neg x_3) \wedge (x_1 \vee \neg x_4) \wedge x_4$, and $\varphi_2 = (x_1 \vee \neg x_3) \wedge (\neg x_1 \vee \neg x_2 \vee \neg x_3) \wedge x_3$. We have $\text{mod}_{\min}(\varphi_1) = \{x_1, x_4\}$ and $\text{mod}_{\min}(\varphi_2) = \{x_1, x_3\}$.

We will compare the behaviour of Θ^{sem} and Θ^{syn} when merging E . Both operators require to determine $H(E) = \langle CL(\varphi_1) \wedge H(V), CL(\varphi_2) \wedge H(V) \rangle$. The translation of φ_1 and φ_2 are formulae with a single model. We use ω_i as a shortcut for $\text{mod}(CL(\varphi_i) \wedge H(V))$, respectively:

$$\begin{aligned} \omega_1 &= \{x_1, x_4, \text{cl}_{x_1 \vee \neg x_2 \vee \neg x_3}, \text{cl}_{x_1 \vee \neg x_4}, \text{cl}_{x_4}\} \\ \omega_2 &= \{x_1, x_3, \text{cl}_{x_1 \vee \neg x_3}, \text{cl}_{\neg x_1 \vee \neg x_2 \vee \neg x_3}, \text{cl}_{x_3}\} \end{aligned}$$

We consider a constraint μ such that $\text{mod}(\mu \wedge H(V)) = \{\omega'_1, \omega'_2\}$ as follows:

$$\begin{aligned} \omega'_1 &= \{x_1, x_3, x_4, \text{cl}_{x_1}, \text{cl}_{x_3}, \text{cl}_{x_4}\} \\ \omega'_2 &= \{\text{cl}_{x_1 \vee \neg x_2 \vee \neg x_3}, \text{cl}_{x_1 \vee \neg x_4}, \text{cl}_{x_1 \vee \neg x_3}, \text{cl}_{\neg x_1 \vee \neg x_2 \vee \neg x_3}\} \end{aligned}$$

The result of the merging, at the level of the propositional translation, is a subset of models of $\mu \wedge H(V)$, which are minimal w.r.t. the distance-based pre-orders corresponding to $\Delta^{(|\text{hcl}(V)|+1, 1), \Sigma}$ (for Θ^{sem}) and $\Delta^{(1, |V|+1), \Sigma}$ (for Θ^{syn}). Table 1 represents the distances between models of $\mu \wedge H(V)$ and the models of the Horn formulae from E , and finally the sum aggregation, expressed with the weights (W_1, W_2) which underly the distances.

$\text{mod}(\mu \wedge H(V))$	ω_1	ω_2	Σ
ω'_1	$W_1 + 4W_2$	$W_1 + 4W_2$	$2W_1 + 8W_2$
ω'_2	$2W_1 + 3W_2$	$2W_1 + 3W_2$	$4W_1 + 6W_2$

Table 1. Distances between interpretation w.r.t. weights (W_1, W_2)

To obtain the result of merging for Θ^{sem} and Θ^{syn} , we need to instantiate the weights in Table 1 with the values corresponding to the actual distances (see Definition 12 and Definition 13). We observe that Θ^{sem} (resp. Θ^{syn}) selects as minimal model ω'_1 (resp. ω'_2). So, $\Theta^{sem}(E) = \{x_1 \wedge x_3 \wedge x_4\}$ and $\Theta^{syn} = \{(x_1 \vee \neg x_2 \vee \neg x_3) \wedge (x_1 \vee \neg x_4) \wedge (x_1 \vee \neg x_3) \wedge (\neg x_1 \vee \neg x_2 \vee \neg x_3)\}$.

5.2 Axiomatization of mod_{\min} -Merging

Before stating the mod_{\min} adaptation of IC-merging postulates, let us introduce some notations. For any E , $\bigwedge \text{mod}_{\min}(E) = \{\omega \mid \forall \varphi \in E, \omega = \text{mod}_{\min}(\varphi)\}$, which means that $\bigwedge \text{mod}_{\min}(E)$ is

empty if all formulae in E do not have the same minimal model. We define the union of profiles E_1 and E_2 as their concatenation. Finally, two Horn profiles E_1 and E_2 are called equivalent, denoted $E_1 \equiv E_2$, if there is a bijective function f from E_1 to E_2 such that $\forall \varphi \in E_1, \text{mod}_{\min}(f(\varphi)) = \text{mod}_{\min}(\varphi)$. Now, a mod_{\min} -merging operator is a mapping from a Horn profile E and a formula μ to a Horn formula $\Theta_\mu(E)$ which satisfies

- ($\Theta 0$) $\text{mod}_{\min}(\Theta_\mu(E)) \subseteq \text{mod}(\mu)$.
- ($\Theta 1$) If $\text{mod}(\mu) \neq \emptyset$, then $\text{mod}_{\min}(\Theta_\mu(E)) \neq \emptyset$.
- ($\Theta 2$) If $\bigwedge \text{mod}_{\min}(E) \cap \text{mod}(\mu) \neq \emptyset$, then $\text{mod}_{\min}(\Theta_\mu(E)) = \bigwedge \text{mod}_{\min}(E) \cap \text{mod}(\mu)$.
- ($\Theta 3$) If $E_1 \equiv E_2$ and $\text{mod}(\mu_1) = \text{mod}(\mu_2)$, then $\Theta_{\mu_1}(E_1) = \Theta_{\mu_2}(E_2)$.
- ($\Theta 4$) If $\{\text{mod}_{\min}(\varphi_1)\} \subseteq \text{mod}(\mu)$ and $\{\text{mod}_{\min}(\varphi_2)\} \subseteq \text{mod}(\mu)$, then $\text{mod}_{\min}(\Theta_\mu(\langle \varphi_1, \varphi_2 \rangle)) \cap \{\text{mod}_{\min}(\varphi_1)\} \neq \emptyset$ iff $\text{mod}_{\min}(\Theta_\mu(\langle \varphi_1, \varphi_2 \rangle)) \cap \{\text{mod}_{\min}(\varphi_2)\} \neq \emptyset$.
- ($\Theta 5$) $\text{mod}_{\min}(\Theta_\mu(E_1)) \cap \text{mod}_{\min}(\Theta_\mu(E_2)) \subseteq \text{mod}_{\min}(\Theta_\mu(E_1 \cup E_2))$.
- ($\Theta 6$) If $\text{mod}_{\min}(\Theta_\mu(E_1)) \cap \text{mod}_{\min}(\Theta_\mu(E_2)) \neq \emptyset$, then $\text{mod}_{\min}(\Theta_\mu(E_1 \cup E_2)) \subseteq \text{mod}_{\min}(\Theta_\mu(E_1)) \cap \text{mod}_{\min}(\Theta_\mu(E_2))$.
- ($\Theta 7$) $\text{mod}_{\min}(\Theta_{\mu_1}(E)) \cap \text{mod}(\mu_2) \subseteq \text{mod}_{\min}(\Theta_{\mu_1 \wedge \mu_2}(E))$.
- ($\Theta 8$) If $\text{mod}_{\min}(\Theta_{\mu_1}(E)) \cap \text{mod}(\mu_2) \neq \emptyset$, then $\text{mod}_{\min}(\Theta_{\mu_1 \wedge \mu_2}(E)) \subseteq \text{mod}_{\min}(\Theta_{\mu_1}(E)) \cap \text{mod}(\mu_2)$.

Similarly to the case of classical propositional belief merging, we can extend the notion of min-faithful assignment, to define pre-orders suitable to define merging operators.

Definition 14. A min-syncretic assignment is a mapping from a Horn profile E to a total pre-order between interpretations \leq_φ such that:

1. $\forall \omega_1 \notin \bigwedge \text{mod}_{\min}(E)$, if $\omega_2 \in \bigwedge \text{mod}_{\min}(E)$ then $\omega_2 <_E \omega_1$;
2. if $E_1 \equiv E_2$, then $\leq_{E_1} = \leq_{E_2}$;
3. if $\omega_1 = \text{mod}_{\min}(\varphi_1)$ and $\omega_2 = \text{mod}_{\min}(\varphi_2)$, then $\omega_1 \simeq_{\langle \varphi_1, \varphi_2 \rangle} \omega_2$;
4. If $\omega_1 \leq_{E_1} \omega_2$ and $\omega_1 \leq_{E_2} \omega_2$, then $\omega_1 \leq_{E_1 \cup E_2} \omega_2$;
5. If $\omega_1 <_{E_1} \omega_2$ and $\omega_1 \leq_{E_2} \omega_2$, then $\omega_1 <_{E_1 \cup E_2} \omega_2$.

Theorem 6. Given a min-syncretic assignment which maps each Horn profile E to a total pre-order between interpretations \leq_E , if the merging operator Θ satisfies

$$\text{mod}_{\min}(\Theta_\mu(E)) = \min(\text{mod}(\mu), \leq_E)$$

then Θ satisfies the postulates ($\Theta 0$)–($\Theta 8$).

Definition 15. Let d be a distance between interpretations. For each Horn profile $E = \langle \varphi_1, \dots, \varphi_n \rangle$, we define a total pre-order between interpretations \leq_E as follows:

$$\omega_1 \leq_E^d \omega_2 \text{ iff } \sum_{\varphi \in E} d(\omega_1, \text{mod}_{\min}(\varphi)) \leq \sum_{\varphi \in E} d(\omega_2, \text{mod}_{\min}(\varphi))$$

The distance-based merging operator Θ^d is defined by

$$\text{mod}_{\min}(\Theta_\mu^d(E)) = \min(\text{mod}(\mu), \leq_E^d)$$

Proposition 7. Every distance-based merging operator satisfies the postulates ($\Theta 0$)–($\Theta 8$).

Proposition 8. The semantic minimal change merging operator Θ^{sem} of Definition 12 satisfies the postulates ($\Theta 0$)–($\Theta 8$), provided the integrity constraint μ is built on V .

6 Reducing the Size of the Encoding

The encoding we propose to express the syntax and semantics of a Horn formula φ is exponentially larger than φ . This stems from the fact that we have to specify, for each possible Horn clause $c \in \text{hcl}(V)$, whether c appears in φ or not. Similarly, we have to specify how c influences the computation of the minimal model. So both $CL(\varphi)$ and $H(V)$ exhibit an exponential blowup. This is the price we have to pay, at least at the conceptual level, for the expressiveness of our approach.

However, for practical applications, we can reduce the size of the encoding. Indeed, it seems reasonable to consider that some clauses are not relevant for the agent (resp. the group of agents) which revises (resp. merges) beliefs, and should not be involved in the revision (resp. merging). For instance, it may be reasonable to assume that only clauses up to a restricted size s , where s is some not too large constant, are relevant. For this reason, $CL(\varphi)$ and $H(V)$ can be adapted to take into account a specific pool of available clauses.

Definition 16. Given $\varphi \in \mathcal{L}_H^{cons}(V)$ and P a set of Horn clauses such that $\text{hcl}(\varphi) \subseteq P$, we define

$$CL_P(\varphi) = \left(\bigwedge_{c \in \text{hcl}(\varphi)} \text{cl}_c \right) \wedge \left(\bigwedge_{c \in P \setminus \text{hcl}(\varphi)} \neg \text{cl}_c \right)$$

Similarly, we define $H_P(V)$ following Definition 2, but re-writing the deduction part to take into account the available clauses:

$$\text{ded}_P^{(x,i)} = \bigvee_{c \in P, c = \neg y_1 \vee \dots \vee \neg y_k \vee x} (\text{cl}_c \wedge y_1^{i-1} \wedge \dots \wedge y_k^{i-1})$$

Given a translation-based revision (resp. merging) operator \star (resp. Θ), we use \star^P (resp. Θ^P) for the corresponding P -based operator.

The minimal pool of clauses P which can be used must contain at least each possible unit clause (to be able to obtain every minimal model), and each clause which is directly involved in the revision (resp. merging) process. These are the clauses c such that cl_c appears explicitly in μ , and the clauses $c \in \text{hcl}(\varphi)$ (resp. $c \in \text{hcl}(\varphi_i)$, for some $\varphi_i \in E$). This minimal pool is polynomially larger than the initial formula (resp. profile). Let us still notice that it is useful to keep a larger pool of clauses in some situations. Especially, with the minimal change of syntax criterion, revision (resp. merging) can lead to a different result if more clauses are available.

Example 7. Let $V = \{x_1, x_2, x_3\}$, and $\varphi = x_1$. We want to revise φ by $\mu = x_3 \wedge \neg \text{cl}_{x_3}$, which means that x_3 must appear in the minimal model, but must not appear as a unit clause. We use the revision operator \star_{syn} . If the revision uses the pool $P = \{\text{cl}_{x_1}, \text{cl}_{x_2}, \text{cl}_{x_3}, \text{cl}_{\neg x_1 \vee x_2}, \text{cl}_{\neg x_2 \vee x_3}\}$, the result of the revision is $\varphi \star_{syn}^P \mu = \{x_1 \wedge (\neg x_1 \vee x_2) \wedge (\neg x_2 \vee x_3)\}$. We observe that two clauses have been added to the initial formula φ . Now, if we consider the pool $P' = P \cup \{\neg x_1 \vee x_3\}$, we notice that it is possible to obtain a better result (w.r.t. minimal change of the syntax): $\varphi \star^{P'} \mu = \{x_1 \wedge (\neg x_1 \vee x_3)\}$.

Of course, all occurrences of $\text{hcl}(V)$ must be replaced by P to define P -based operators (e.g. in Definition 4 to define $d_{(W_1, W_2)}$ on the adequate set of variables).

7 Discussion

Related Work As far as we know, this is the first work which considers belief revision or merging in a situation where the relevant information is the minimal model of a formula. In particular, previous

works on Horn revision and merging consider the whole set of classical models. These approaches are classified along two lines. First, [7, 8] proposes to refine existing revision operators to guarantee that the result belongs to a given fragment. These refined operators coincide with the original ones in the case when the result of the original operator belongs to the target fragment. The second research direction is a modification of the revision and merging postulates [10, 19] and the definition of new operators corresponding to these postulates to ensure that the result belongs to the Horn fragment.

Also belief change for logic programs is related to our work, since sets of Horn clauses under minimal model semantics are a special case of logic programs under stable semantics. In fact, there is a huge body of work which addresses (usually a syntactic approach for) update, see e.g. [2, 15]. Hereby, dedicated properties different to AGM-style postulates have been proposed; hence, this research branch significantly differs from our approach. However, there is also work which is using the notion of SE-models [22, 27] as the objects revision [11, 25] or merging [12] operators are defined on. SE-models capture classical models of programs together with models of the respective program reducts. For the special case of Horn programs, reduct models and program models coincide. Thus, SE-models amount to classical models. Hence, for Horn programs these approaches are equivalent to the ones for Horn revision and merging based on classical models [10, 19]. Given the wide range of work in revision or update of logic programs, we are not aware of any approach that allows for simultaneous control (via the revision formula) whether minimality of change primarily relates to the syntax or to semantics of the program, in the way our method does.

The first work on translation-based revision is [17] which uses translations in classical logic to revise theories from other settings, like modal logic K , Łukasiewicz's finitely many-valued logic \mathcal{L}_n , algebraic logic and Belnap's four-valued logic. These translations allow to revise a theory by a formula from the same formalism, but do not provide any means to revise the information about the syntax of formulae. The work described in [6] is closer to our contribution. Here, argumentation frameworks are translated into propositional formulae which can be revised by a piece of information about the semantics (arguments statuses) and the syntax (attack relation) of the framework.

Conclusion and Future Research In this paper, we have defined an original translation-based approach to revise and merge consistent Horn formulae, in the special scenario of agents reasoning based on the minimal model of the formulae. We have defined specific revision and merging operators, proposed rationality postulates for both operations, and we have shown that some of our operators satisfy the postulates. Relaxing this restriction to consistent formulae is not straightforward, and would unnecessarily complexify the whole definition of the process.

With our adaptation of KM revision postulates and IC-merging postulates, we have proved that some of our operators are related to purely model-based revision and merging. In future work, we want to investigate the relation between our translation-based approach and syntax-based revision and merging. Indeed, we think that this kind of translation-based approach is a perfect way to explore the "middle ground" between syntactic and model-based approaches to belief change, combining the importance of both aspects.

There are various interesting open questions which we want to address in future work. First of all, our approach is nondeterministic in the sense that it may produce a collection of Horn formulae, each of them representing an acceptable result of revision or merging. This

kind of nondeterminism is quite common, especially in nonmonotonic reasoning where, say, a default theory can produce multiple extensions, or a logic program may have multiple stable models. Nevertheless, it is interesting to further investigate this issue. On one hand, conditions under which a unique Horn formula is generated are of interest. This is the case in particular when propositional operators are used, at the level of the encoding, which always yield a formula with a single model. On the other hand, we want to study selection functions for determining a unique outcome. Such functions may be based on further criteria, like additional preference information, or on some kind of a tie-break rule. A selection function must be used whenever a specific application requires a single Horn formula as the result of the process.

Finally, we aim for generalizing our logical encoding to take all models of a Horn formula into account, not only the minimal one. Such a generalization might pave the way to apply our translation-based method to other formalisms of interest, such as full propositional logic and logic programs.

Acknowledgements. G. Brewka was supported by DFG research unit FOR 1513 (HYBRIS). J.-G. Mailly and S. Woltran were supported by the Austrian Science Fund (FWF) under grant P25521.

A Proofs

Some proofs are omitted due to space reasons. In particular, proofs about translation-based merging are similar to proofs about translation based revision, with special care to the aggregation function.

Proof of Proposition 2. First, let us notice that if two interpretations ω_1, ω_2 coincide on the variables from $V \cup \{\text{cl}_c \mid c \in \text{hcl}(V)\}$, then $\omega_1 = \omega_2$, since the variables x^i can be deduced from the other variables. We call M the set of models of $CL(\varphi) \wedge H(V)$.

$\leq_{\varphi}^{(W_1, W_2)}$ denotes the pre-order defined by $\omega_1 \leq_{\varphi}^{(W_1, W_2)} \omega_2$ iff $\min_{\omega' \in M} d_{(W_1, W_2)}(\omega_1, \omega') \leq \min_{\omega' \in M} d_{(W_1, W_2)}(\omega_2, \omega')$.

Let $\omega_1 \in M$. Obviously, $\min_{\omega' \in M} d_{(W_1, W_2)}(\omega_1, \omega') = 0$. Now, given another interpretation ω_2 , $\min_{\omega' \in M} d_{(W_1, W_2)}(\omega_2, \omega') = 0$ holds iff $\omega_2 \in M$. In this case, $\omega_1 \simeq_{\varphi}^{(W_1, W_2)} \omega_2$, otherwise $\omega_1 <_{\varphi}^{(W_1, W_2)} \omega_2$. If we consider a formula $\varphi' \equiv \varphi$, the distance do not change (since the set of models are identical), and the pre-order $\leq_{\varphi'}^{(W_1, W_2)}$ is equal to $\leq_{\varphi}^{(W_1, W_2)}$.

The properties of faithful assignments are satisfied. This means that the revision operator induced by $\leq_{\varphi}^{(W_1, W_2)}$ satisfies the KM rationality postulates. \square

Proof of Theorem 3. Let \star be a revision operator based on a min-faithful assignment. \star satisfies $(\star 1)$ by definition.

$\{\text{mod}_{\min}(\varphi)\} \cap \text{mod}(\mu) \neq \emptyset$ is equivalent to $\text{mod}_{\min}(\varphi) \in \text{mod}(\mu)$. Since $\text{mod}_{\min}(\varphi)$ is the minimal element w.r.t \leq_{φ} , it is obviously the minimal model of μ w.r.t \leq_{φ} , so $(\star 2)$ is satisfied.

If $\text{mod}(\mu) \neq \emptyset$, then $\min(\text{mod}(\mu), \leq_{\varphi}) \neq \emptyset$ is ensured since \leq_{φ} is a total relation, so $(\star 3)$ is satisfied.

$(\star 4)$ follows from the definition of \leq_{φ} and \star . Indeed, if $\text{mod}_{\min}(\varphi_1) = \text{mod}_{\min}(\varphi_2)$ and $\text{mod}(\mu_1) = \text{mod}(\mu_2)$,

$$\begin{aligned} \text{mod}_{\min}(\varphi_1 \star \mu_1) &= \min(\text{mod}(\mu_1), \leq_{\varphi_1}) \\ &= \min(\text{mod}(\mu_2), \leq_{\varphi_1}) \text{ since } \mu_1 \equiv \mu_2 \\ &= \min(\text{mod}(\mu_2), \leq_{\varphi_2}) \text{ from Def.7} \\ &= \text{mod}_{\min}(\varphi_2 \star \mu_2) \end{aligned}$$

If $\text{mod}_{\min}(\varphi \circ \mu_1) \cap \text{mod}(\mu_2) = \emptyset$, then the last postulates are satisfied. So now we suppose that $\text{mod}_{\min}(\varphi \circ \mu_1) \cap \text{mod}(\mu_2) \neq \emptyset$.

We have $\text{mod}_{\min}(\varphi \circ \mu_1) \cap \text{mod}(\mu_2) = \min(\text{mod}(\mu_1), \leq_{\varphi}) \cap \text{mod}(\mu_2)$, and $\text{mod}_{\min}(\varphi \circ (\mu_1 \wedge \mu_2)) = \min(\text{mod}(\mu_1 \wedge \mu_2), \leq_{\varphi})$.

Using *reductio ad absurdum*, we suppose that there exists some interpretation ω_1 such that $\omega_1 \in \min(\text{mod}(\mu_1), \leq_{\varphi}) \cap \text{mod}(\mu_2)$ and $\omega_1 \notin \min(\text{mod}(\mu_1 \wedge \mu_2), \leq_{\varphi})$. From the first part, we know that $\omega_1 \in \text{mod}(\mu_1)$ and $\omega_1 \in \text{mod}(\mu_2)$, i.e. $\omega_1 \in \text{mod}(\mu_1 \wedge \mu_2)$. Now we deduce $\exists \omega_2 \in \text{mod}(\mu_1 \wedge \mu_2)$ such that $\omega_2 <_{\varphi} \omega_1$; but $\omega_2 \in \text{mod}(\mu_1)$, so $\omega_1 \notin \min(\text{mod}(\mu_1), \leq_{\varphi})$. This is a contradiction so \star satisfies $(\star 5)$.

To prove the opposite inclusion, let us also reason with *reduction ad absurdum*. We suppose that there is an interpretation ω_1 such that $\omega_1 \in \min(\text{mod}(\mu_1 \wedge \mu_2), \leq_{\varphi})$ and $\omega_1 \notin \min(\text{mod}(\mu_1), \leq_{\varphi}) \cap \text{mod}(\mu_2)$. From the first part, $\omega_1 \in \text{mod}(\mu_1 \wedge \mu_2)$, stated otherwise: $\omega_1 \in \text{mod}(\mu_1)$ and $\omega_1 \in \text{mod}(\mu_2)$. Since we have suppose that $\text{mod}_{\min}(\varphi \circ \mu_1) \cap \text{mod}(\mu_2) \neq \emptyset$, there is an interpretation ω_2 such that $\omega_2 \in \min(\text{mod}(\mu_1), \leq_{\varphi}) \cap \text{mod}(\mu_2)$. From these, we deduce $\omega_2 <_{\varphi} \omega_1$ and $\omega_2 \in \text{mod}(\mu_1) \cap \text{mod}(\mu_2) = \text{mod}(\mu_1 \wedge \mu_2)$, and so $\omega_1 \notin \min(\text{mod}(\mu_1 \wedge \mu_2), \leq_{\varphi})$. This is a contradiction, so \star satisfies $(\star 6)$. \square

Proof of Proposition 4. We just need to prove that mapping a Horn formula to a distance-based pre-order defines a min-faithful assignment. Let d be a distance between interpretations. $\forall \omega \neq \text{mod}_{\min}(\varphi)$, $d(\text{mod}_{\min}(\varphi), \text{mod}_{\min}(\varphi)) = 0$ is strictly less than $d(\omega, \text{mod}_{\min}(\varphi))$, so $\text{mod}_{\min}(\varphi) <_{\varphi}^d \omega$. $\forall \omega_1, \omega_2$, $\omega_1 \leq_{\varphi_1}^d \omega_2$ iff $d(\omega_1, \text{mod}_{\min}(\varphi_1)) \leq d(\omega_2, \text{mod}_{\min}(\varphi_1))$. Under the assumption $\text{mod}_{\min}(\varphi_1) = \text{mod}_{\min}(\varphi_2)$, this is equivalent to $d(\omega_1, \text{mod}_{\min}(\varphi_2)) \leq d(\omega_2, \text{mod}_{\min}(\varphi_2))$, and so $\omega_1 \leq_{\varphi_2}^d \omega_2$. This means that the pre-orders \leq_{φ_1} and \leq_{φ_2} are equals. \square

Proof of Proposition 5. We want to prove that the semantic minimal change revision operator is a particular distance based revision operator. $\star_{(W_1, 1)}$, with $W_1 = |\text{hcl}(V)| + 1$, is built from a distance between interpretations on the set of variables $V' = V \cup \{\text{cl}_c \mid c \in \text{hcl}(V)\}$. To be able to apply Proposition 4, we need to reformulate the definition of the operator to represent it with a distance between propositional interpretations on the set of variables V .

First of all, we observe from the definition of the translation-based revision operators that $\text{mod}_{\min}(\text{dec}(\varphi \star_{(W_1, 1)} \mu)) \subseteq \text{mod}(\mu)$. We need to identify a distance d on V such that $\text{mod}_{\min}(\text{dec}(\varphi \star_{(W_1, 1)} \mu)) = \min(\text{mod}(\mu), \leq_d)$. In our case, the Hamming distance on V proves enough to obtain the result.

Given $\omega_1 \in \min(\text{mod}(\mu \wedge H(V)), \leq_{\varphi}^{d_{(W_1, 1)}})$, the distance $d_{(W_1, 1)}(\omega_1, \text{mod}(CL(\varphi) \wedge H(V)))$ is minimal. In particular, since $CL(\varphi) \wedge H(V)$ has a single model ω_{φ} , $d_{(W_1, 1)}(\omega_1, \omega_{\varphi})$ is minimal. Stated otherwise, $W_1 d_H^V(\omega_1, \omega_{\varphi}) + d_H^{\text{Syn}}(\omega_1, \omega_{\varphi})$ is the minimal distance between ω_{φ} and another interpretation. Let us suppose that the x -part is not minimal, i.e. there exists ω_2 such that $W_1 d_H^V(\omega_2, \omega_{\varphi}) < W_1 d_H^V(\omega_1, \omega_{\varphi})$. We consider the extreme case when the cl_c -part is nul for ω_1 , i.e. $d_H^{\text{Syn}}(\omega_1, \omega_{\varphi}) = 0$, and it is maximal for ω_2 , i.e. $d_H^{\text{Syn}}(\omega_2, \omega_{\varphi}) = |\text{hcl}(V)|$. Then we obtain

$$\begin{aligned} W_1 d_H^V(\omega_2, \omega_{\varphi}) + d_H^{\text{Syn}}(\omega_2, \omega_{\varphi}) \\ < W_1 d_H^V(\omega_1, \omega_{\varphi}) + d_H^{\text{Syn}}(\omega_1, \omega_{\varphi}) \end{aligned}$$

because $W_1 = |\text{hcl}(V)| + 1$. This means that $d_{(W_1, 1)}(\omega_2, \omega_{\varphi}) < d_{(W_1, 1)}(\omega_1, \omega_{\varphi})$, which is a contradiction. This means that for each $\omega_1 \in \min(\text{mod}(\mu \wedge H(V)), \leq_{\varphi}^{d_{(W_1, 1)}})$, and for each ω_2 , $W_1 d_H^V(\omega_1, \omega_{\varphi}) \leq W_1 d_H^V(\omega_2, \omega_{\varphi})$, which is equivalent to $d_H^V(\omega_1, \omega_{\varphi}) \leq d_H^V(\omega_2, \omega_{\varphi})$. We notice that this means that the x -part of ω_1 is a minimal interpretation on V with respect to the Hamming distance. \square

REFERENCES

- [1] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson, 'On the logic of theory change: Partial meet contraction and revision functions', *J. Symb. Log.*, **50**(2), 510–530, (1985).
- [2] José Júlio Alferes, João Alexandre Leite, Luís Moniz Pereira, Halina Przymusinska, and Teodor C. Przymusinski, 'Dynamic updates of non-monotonic knowledge bases', *Journal of Logic Programming*, **45**(1–3), 43–70, (2000).
- [3] Grigoris Antoniou and Mary-Anne Williams, 'Revising default theories', in *Proceedings of the Tenth IEEE International Conference on Tools with Artificial Intelligence (ICTAI'98)*, pp. 423–430, (1998).
- [4] Ringo Baumann and Gerhard Brewka, 'AGM meets abstract argumentation: Expansion and revision for Dung frameworks', in *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI'15)*, pp. 2734–2740, (2015).
- [5] Sylvie Coste-Marquis, Sébastien Konieczny, Jean-Guy Mailly, and Pierre Marquis, 'On the revision of argumentation systems: Minimal change of arguments statuses', in *Proceedings 14th International Conference on Principles of Knowledge Representation and Reasoning (KR'14)*, pp. 52–61, (2014).
- [6] Sylvie Coste-Marquis, Sébastien Konieczny, Jean-Guy Mailly, and Pierre Marquis, 'A translation-based approach for revision of argumentation frameworks', in *Proceedings of the Fourteenth European Conference on Logics in Artificial Intelligence (JELIA'14)*, pp. 397–411, (2014).
- [7] Nadia Creignou, Odile Papini, Reinhard Pichler, and Stefan Woltran, 'Belief revision within fragments of propositional logic', *J. Comput. Syst. Sci.*, **80**(2), 427–449, (2014).
- [8] Nadia Creignou, Odile Papini, Stefan Rümmele, and Stefan Woltran, 'Belief merging within fragments of propositional logic', in *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI'14)*, pp. 231–236, (2014).
- [9] Nadia Creignou, Reinhard Pichler, and Stefan Woltran, 'Do hard SAT-related reasoning tasks become easier in the Krom fragment?', in *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI'13)*, pp. 824–831, (2013).
- [10] James P. Delgrande and Pavlos Peppas, 'Belief revision in Horn theories', *Artif. Intell.*, **218**, 1–22, (2015).
- [11] James P. Delgrande, Pavlos Peppas, and Stefan Woltran, 'AGM-style belief revision of logic programs under answer set semantics', in *Proceedings of the Twelfth International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'13)*, pp. 264–276, (2013).
- [12] James P. Delgrande, Torsten Schaub, Hans Tompits, and Stefan Woltran, 'A model-theoretic approach to belief change in answer set programming', *ACM Trans. Comput. Log.*, **14**(2), 14, (2013).
- [13] Martin Diller, Adrian Haret, Thomas Linsbichler, Stefan Rümmele, and Stefan Woltran, 'An extension-based approach to belief revision in abstract argumentation', in *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI'15)*, pp. 2926–2932, (2015).
- [14] Phan Minh Dung, 'On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n-person games', *Artif. Intell.*, **77**(2), 321–357, (1995).
- [15] Thomas Eiter, Michael Fink, Giuliana Sabbatini, and Hans Tompits, 'On properties of update sequences based on causal rejection', *TPLP*, **2**(6), 711–767, (2002).
- [16] Dov M. Gabbay, Odinaldo Rodrigues, and Alessandra Russo, *Revision, Acceptability and Context - Theoretical and Algorithmic Aspects*, Cognitive Technologies, Springer, 2010.
- [17] Dov M. Gabbay, Odinaldo Rodrigues, and Alessandra Russo, 'Revision by translation', in *Revision, Acceptability and Context*, 223–270, Springer, (2010).
- [18] Peter Gärdenfors, 'Belief revision and nonmonotonic logic: Two sides of the same coin?', in *Proceedings of the Ninth European Conference on Artificial Intelligence (ECAI'90)*, pp. 768–773, (1990).
- [19] Adrian Haret, Stefan Rümmele, and Stefan Woltran, 'Merging in the Horn fragment', in *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI'15)*, pp. 3041–3047, (2015).
- [20] Hirofumi Katsuno and Alberto O. Mendelzon, 'Propositional knowledge base revision and minimal change', *Artificial Intelligence*, **52**, 263–294, (1991).
- [21] Sébastien Konieczny and Ramón Pino Pérez, 'Merging information under constraints: A logical framework', *J. Log. Comput.*, **12**(5), 773–808, (2002).
- [22] Vladimir Lifschitz, David Pearce, and Agustin Valverde, 'Strongly equivalent logic programs', *ACM Trans. Comput. Log.*, **2**(4), 526–541, (2001).
- [23] Jean-Guy Mailly, *Dynamics of Argumentation Frameworks*, Ph.D. dissertation, Université d'Artois, 2015.
- [24] Guilin Qi and Jianfeng Du, 'Model-based revision operators for terminologies in description logics', in *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI'09)*, pp. 891–897, (2009).
- [25] Nicolas Schwind and Katsumi Inoue, 'Characterization of logic program revision as an extension of propositional revision', *TPLP*, **16**(1), 111–138, (2016).
- [26] Martin Slota and João Leite, 'The rise and fall of semantic rule updates based on se-models', *TPLP*, **14**(6), 869–907, (2014).
- [27] Hudson Turner, 'Strong equivalence made easy: nested expressions and weight constraints', *TPLP*, **3**(4–5), 609–622, (2003).