Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

Merging in the Horn Fragment*

Adrian Haret, Stefan Rümmele, Stefan Woltran
{haret, ruemmele, woltran}@dbai.tuwien.ac.at
Institute of Information Systems
Vienna University of Technology, Austria

Abstract

Belief merging is a central operation within the field of belief change and addresses the problem of combining multiple, possibly mutually inconsistent knowledge bases into a single, consistent one. A current research trend in belief change is concerned with tailored representation theorems for fragments of logic, in particular Horn logic. Hereby, the goal is to guarantee that the result of the change operations stays within the fragment under consideration. While several such results have been obtained for Horn revision and Horn contraction, merging of Horn theories has been neglected so far. In this paper, we provide a novel representation theorem for Horn merging by strengthening the standard merging postulates. Moreover, we present a concrete Horn merging operator satisfying all postulates.

1 Introduction

Belief merging uses a logical approach to study how information coming from multiple, possibly mutually inconsistent knowledge bases should be combined to form a single, consistent knowledge base. Merging shares a common methodology with other belief change operators, such as revision [Alchourrón *et al.*, 1985; Katsuno and Mendelzon, 1992], contraction [Alchourrón *et al.*, 1985] and update [Katsuno and Mendelzon, 1991]. Part of the methodology is the formulation of postulates, which any rational operator should satisfy. For merging, the IC-merging postulates [Konieczny and Pino Pérez, 2002; 2011] are commonly used. In a further step, a representation result is usually derived: this shows that all (merging) operators satisfying the postulates can be characterized using rankings on the possible worlds described by the underlying language, which is typically taken to be full propositional logic.

Recently, the restriction of belief change formalisms to different fragments of propositional logic has become a vivid research branch. There are pragmatic reasons for focusing on fragments, especially Horn logic. Firstly, Horn clauses are a natural way of formulating basic facts and rules, and thus are useful to encode expert knowledge. Second, Horn logic affords very efficient algorithms. Thus the computational cost of reasoning in this fragment is comparatively low.

While revision [Delgrande and Peppas, 2015; Van De Putte, 2013; Zhuang *et al.*, 2013] and contraction [Booth *et al.*, 2011; Delgrande and Wassermann, 2013; Zhuang and Pagnucco, 2012] have received a lot of attention in this direction, belief merging has yet remained unexplored, with the notable exception of [Creignou *et al.*, 2014]. We aim to fill this gap and investigate the problem of merging in the Horn fragment of propositional logic. We find that restricting the underlying language poses a series of non-trivial challenges, as representation results which work for full propositional logic break down in the Horn case.

Firstly, we find that we cannot rely on the same types of rankings as the ones used for merging in the case of full propositional logic. The reason is that such rankings lead to outputs that can not be expressed as Horn formulas. We fix this problem by adding the restriction of *Horn compliance*: this narrows down the notion of ranking in a way that is coherent with the semantics of Horn formulas. Since standard merging operators are found not to be Horn compliant (hence useless for our needs), we also give a concrete operator that exhibits this property. This is remarkable, since previous research [Creignou *et al.*, 2014] only resulted in Horn merging operators that do not satisfy all postulates. Secondly, Horn merging operators that satisfy the standard postulates turn out to represent more rankings than was expected, some of which are undesirable. We interpret this as an inadequacy of the standard postulates to capture the intended intuitive behaviour of a merging operator. Hence, we propose an alternative formulation of some key postulates, which allows us to derive an appealing representation result for the case of Horn merging. Our approach here is inspired by existing work on Horn revision [Delgrande and Peppas, 2015], though we go significantly beyond it to tackle the problems posed by merging.

The rest of the paper is organized as follows. In Section 2 we introduce the background to merging. In Section 3 we argue that standard model-based merging operators are inappropriate for Horn merging and introduce the property of Horn compliance. In Section 4 we argue that a subset of the IC-merging postulates should be replaced by a strengthened version, and introduce a representation result using the strengthened postulates. Finally, in Section 5 we describe a

*Supported by the Austrian Science Fund (FWF) under grants P25518 and P25521.

concrete Horn merging operator satisfying all postulates. Due to lack of space, we do not include here the proofs of the claims found in the text. These can be found in the full version [Haret *et al.*, 2015].

2 Preliminaries

Propositional logic. We work with the language $\mathcal{L}$ of propositional logic over a fixed *alphabet* $\mathcal{P} = \{p_1, \ldots, p_n\}$ of propositional atoms. We use standard connectives $\vee$, $\wedge$, $\neg$ and the logical constants $\top$ and $\bot$. A literal is an atom or a negated atom. A clause is a disjunction of literals. A clause is called *Horn* if at most one of its literals is positive. The Horn fragment $\mathcal{L}_H \subset \mathcal{L}$ is the set of all formulas in $\mathcal{L}$ that are conjunctions of Horn clauses. An interpretation is a set $w \subseteq \mathcal{P}$ of atoms. The set of all interpretations is denoted by $\mathcal{W}$. We will typically represent an interpretation by its corresponding bit-vector of length $|\mathcal{P}|$. As an example, if $|\mathcal{P}| = 3$, then 101 is the interpretation $\{p_1, p_3\}$. A pre-order $\leq$ on $\mathcal{W}$ is a reflexive, transitive binary relation on $\mathcal{W}$. If $w_1, w_2 \in \mathcal{W}$, then $w_1 < w_2$ denotes the strict part of $\leq$, i.e., $w_1 \leq w_2$ but $w_2 \not\leq w_1$. We write $w_1 \approx w_2$ to abbreviate $w_1 \leq w_2$ and $w_2 \leq w_1$. If $\mathcal{M}$ is a set of interpretations, then the set of *minimal elements of* $\mathcal{M}$ *with respect to* $\leq$ is defined as $min_{\leq}\mathcal{M} = \{w_1 \in \mathcal{M} \mid \nexists w_2 \in \mathcal{M} \text{ s.t. } w_2 < w_1\}$. If interpretation w satisfies formula φ, we call w *a model of* φ. We denote the set of models of φ by $[\varphi]$. Given a set $\mathcal{M}$ of interpretations, we define $\mathrm{Cl}_\cap(\mathcal{M})$, the closure of $\mathcal{M}$ under intersection, as the smallest superset of $\mathcal{M}$ that is closed under $\cap$, i.e., if $w_1, w_2 \in \mathrm{Cl}_\cap(\mathcal{M})$ then also $w_1 \cap w_2 \in \mathrm{Cl}_\cap(\mathcal{M})$. We recall here a classic result concerning Horn formulas and their models (see e.g. [Schaefer, 1978]).

Proposition 1. *A set of interpretations $\mathcal{M}$ is the set of models of a Horn formula φ if and only if $\mathcal{M} = Cl_\cap(\mathcal{M})$.*

A formula is called *complete* if it has exactly one model. If w_i is an interpretation, we sometimes write σ_{w_i} or σ_i to denote the complete formula that has w_i as a model. If σ_i and σ_j are complete formulas, then $\sigma_{i,j}$ is a formula such that $[\sigma_{i,j}] = \{w_i, w_j\}$. If we are working in the Horn fragment, we take $\sigma_{i,j}$ to be such that $[\sigma_{i,j}] = Cl_\cap(\{w_i, w_j\})$.

Belief Merging. A *knowledge base* is a finite set of propositional formulas. A *profile* is a non-empty finite multi-set $E = \{K_1, \ldots, K_n\}$ of consistent knowledge bases. *Horn knowledge bases* and *Horn profiles* contain only Horn formulas and Horn knowledge bases, respectively. The sets of all knowledge bases, Horn knowledge bases, profiles and Horn profiles are denoted by $\mathcal{K}$, $\mathcal{K}_H$, $\mathcal{E}$ and $\mathcal{E}_H$, respectively. If E_1 and E_2 are profiles, then $E_1 \sqcup E_2$ is the multi-set union of E_1 and E_2. Interpretation w is a model of knowledge base K if it is a model of every formula in K. Interpretation w is a model of profile E if it is a model of every $K \in E$. We denote by $[K]$ and $[E]$ the set of models of K and E, respectively. We write $\bigwedge E$ for $\bigwedge_{K \in E} \bigwedge_{\varphi \in K} \varphi$. This reduces a profile to a single propositional formula. Clearly, $[\bigwedge E] = [E]$.

Profiles E_1 and E_2 are *equivalent*, written $E_1 \equiv E_2$, if there exists a bijection $f\colon E_1 \to E_2$ such that for any $K \in E_1$ we have $[K] = [f(K)]$.

A merging operator is a function $\Delta\colon \mathcal{E} \times \mathcal{L} \to \mathcal{K}$. It maps a profile E and a formula μ, typically referred to as constraint, onto a knowledge base. We write $\Delta_\mu(E)$ instead of $\Delta(E, \mu)$. As is common in the belief change literature, logical postulates are employed to set out properties which any merging operator Δ should satisfy. An operator satisfying the following postulates is called IC-merging operator [Konieczny and Pino Pérez, 2002; 2011]:

(IC_0) $\Delta_\mu(E) \models \mu$.

(IC_1) If μ is consistent, then $\Delta_\mu(E)$ is consistent.

(IC_2) If $\bigwedge E$ is consistent with μ, then $\Delta_\mu(E) \equiv \bigwedge E \wedge \mu$.

(IC_3) If $E_1 \equiv E_2$ and $\mu_1 \equiv \mu_2$, then $\Delta_{\mu_1}(E_1) \equiv \Delta_{\mu_2}(E_2)$.

(IC_4) If $K_1 \models \mu$ and $K_2 \models \mu$, then $\Delta_\mu(\{K_1, K_2\}) \wedge K_1$ is consistent iff $\Delta_\mu(\{K_1, K_2\}) \wedge K_2$ is consistent.

(IC_5) $\Delta_\mu(E_1) \wedge \Delta_\mu(E_2) \models \Delta_\mu(E_1 \sqcup E_2)$.

(IC_6) If $\Delta_\mu(E_1) \wedge \Delta_\mu(E_2)$ is consistent, then $\Delta_\mu(E_1 \sqcup E_2) \models \Delta_\mu(E_1) \wedge \Delta_\mu(E_2)$.

(IC_7) $\Delta_{\mu_1}(E) \wedge \mu_2 \models \Delta_{\mu_1 \wedge \mu_2}(E)$.

(IC_8) If $\Delta_{\mu_1}(E) \wedge \mu_2$ is consistent, then $\Delta_{\mu_1 \wedge \mu_2}(E) \models \Delta_{\mu_1}(E) \wedge \mu_2$.

It turns out that we can use a particular type of rankings on interpretations to compute the models of a merging operator.

Definition 1. A *syncretic assignment* is a function mapping each $E \in \mathcal{E}$ to a pre-order $\leq_E$ on $\mathcal{W}$ such that, for any $E, E_1, E_2 \in \mathcal{E}$, $K_1, K_2 \in \mathcal{K}$ and $w_1, w_2 \in \mathcal{W}$ the following conditions hold:

(s_1) If $w_1 \in [E]$ and $w_2 \in [E]$, then $w_1 \approx_E w_2$.

(s_2) If $w_1 \in [E]$ and $w_2 \notin [E]$, then $w_1 <_E w_2$.

(s_3) If $E_1 \equiv E_2$, then $\leq_{E_1} = \leq_{E_2}$.

(s_4) If $w_1 \in [K_1]$, then there is $w_2 \in [K_2]$ such that $w_2 \leq_{\{K_1, K_2\}} w_1$.

(s_5) If $w_1 \leq_{E_1} w_2$ and $w_1 \leq_{E_2} w_2$, then $w_1 \leq_{E_1 \sqcup E_2} w_2$.

(s_6) If $w_1 \leq_{E_1} w_2$ and $w_1 <_{E_2} w_2$, then $w_1 <_{E_1 \sqcup E_2} w_2$.

We define syncretic assignments in a way that allows the pre-orders $\leq_E$ to be partial, as we will make use of partial pre-orders in our own results on Horn operators (Theorems 3 and 4). In the context of full propositional logic, however, the classical result below characterizes all IC-merging operators in terms of syncretic assignments with total pre-orders.

Theorem 1 ([Konieczny and Pino Pérez, 2002; 2011]). *A merging operator Δ is an IC-merging operator if and only if there exists a syncretic assignment mapping each $E \in \mathcal{E}$ to a total pre-order $\leq_E$ such that $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$, for any $\mu \in \mathcal{L}$.*

When this equation holds we will say that the assignment represents the operator.

It is useful to think of a profile $E = \{K_1, \ldots, K_n\}$ as a multi-set of agents, represented by their sets of beliefs K_i. Each agent is equipped with a pre-order $\leq_{K_i}$ on $\mathcal{W}$ which can be thought of as the way in which the agent ranks possible worlds in terms of their plausibility. Merging is then the task of finding a common ranking that approximates, as best as

	$[K]$	Σ	$GMAX$
00	2	2	(2)
01	1	1	(1)
10	1	1	(1)
11	0	0	(0)

Table 1: $\Delta_{\mu}^{d_H,\Sigma}(E_1)$ and $\Delta_{\mu}^{d_H,GMAX}(E_1)$ do not stay in the Horn fragment.

	$[K_1]$	$[K_2]$	Σ	$GMAX$
00	1	1	2	(1,1)
01	1	0	1	(1,0)
10	0	1	1	(1,0)
11	1	1	2	(1,1)

Table 2: $\Delta_{\mu}^{d_D,\Sigma}(E_2)$ and $\Delta_{\mu}^{d_D,GMAX}(E_2)$ do not stay in the Horn fragment.

possible, the individual rankings. Proposition 1 tells us that if this process is done using syncretic assignments, we are in agreement with postulates IC_0 – IC_8.

Two parts need to be filled out to get a concrete merging operator: how to compute the individual rankings and how to aggregate them. For the first part, the common approach in the literature is to use some notion of distance between interpretations, such as Hamming distance d_H or the drastic distance d_D. The minimal distance between interpretations and models of K_i is used to construct $\leq_{K_i}$. For the second part, common functions used to aggregate the distances are the sum Σ or $GMAX$, giving us operators $\Delta^{d_H,\Sigma}$, $\Delta^{d_H,GMAX}$, $\Delta^{d_D,\Sigma}$, $\Delta^{d_H,GMAX}$. The reader is referred to [Konieczny and Pino Pérez, 2002; 2011] for more details.

3 Restricting assignments

A Horn merging operator is a function $\Delta\colon \mathcal{E}_H \times \mathcal{L}_H \rightarrow \mathcal{K}_H$. Our aim is to characterize the class of such operators in the manner of Proposition 1 and to exhibit a concrete operator. The first problem occurs when we apply standard merging operators to Horn profiles and formulas: their output cannot always be represented by a Horn formula.

Examples 1 and 2 show how standard merging operators fail in the Horn fragment. In both cases we choose Horn profiles over the 2-letter alphabet and construct rankings using the Hamming distance and the drastic distance. We then aggregate the rankings with Σ and $GMAX$. The rankings and the result of their aggregation are shown in Tables 1 and 2. Each row displays one possible interpretation over $\{p_1, p_2\}$ (denoted in the first column). The second column displays the minimal distance of the interpretation to any model of K. The third and fourth column show the aggregation of the distances (in our case only one distance) according to the Σ as well as the $GMAX$ function. The output of the merging operator is the set of those interpretations that are models of μ (marked grey) and have the smallest aggregated value.

Example 1. Take $K = \{p_1,\ p_2\}$, $E_1 = \{K\}$ and $\mu = \neg p_1 \vee \neg p_2$ (all of them Horn). We compute $\Delta_{\mu}^{d_H,\Sigma}(E_1)$ and $\Delta_{\mu}^{d_H,GMAX}(E_1)$, keeping in mind that $[K] = \{11\}$ and $[\mu] = \{00, 10, 01\}$. Table 1 displays how the output of these two mergings is computed.

We get $[\Delta_{\mu}^{d_H,\Sigma}(E_1)] = min[\mu] = \{10, 01\}$, and the same result is obtained for $\Delta_{\mu}^{d_H,GMAX}(E_1)$. It holds, then, that $\Delta_{\mu}^{d_H,\Sigma}(E_1)$ and $\Delta_{\mu}^{d_H,GMAX}(E_1)$ cannot be expressed as Horn formulas.

Example 2. Take $K_1 = \{p_1,\ \neg p_2\}$, $K_2 = \{\neg p_1,\ p_2\}$, $E_2 = \{K_1,\ K_2\}$ and $\mu = \neg p_1 \vee \neg p_2$: We get (see Table 2) $[\Delta_{\mu}^{d_D,\Sigma}(E_2)] = min[\mu] = \{10, 01\}$, and the same result is obtained for $\Delta_{\mu}^{d_D,GMAX}(E_2)$. Again, $\Delta_{\mu}^{d_D,\Sigma}(E_2)$ and $\Delta_{\mu}^{d_D,GMAX}(E_2)$ cannot be represented by Horn formulas.

As we see in Examples 1 and 2, standard distance and aggregate functions are not adequate for the Horn fragment. Here we adopt a solution suggested by [Delgrande and Peppas, 2015], which is to impose an extra condition on the pre-orders.

Definition 2. A pre-order $\leq$ is *Horn compliant* if for any $\mu \in \mathcal{L}_H$, $min_{\leq}[\mu]$ can be represented by a Horn formula.

Example 3. The computed pre-orders for E_1 and E_2 in Examples 1 and 2 are not Horn compliant, as we get that $min[\mu] = \{01, 10\}$ in both cases.

Adding Horn compliance makes it possible to define a merging operator for the Horn fragment, and this gives us one direction of a representation theorem.

Theorem 2. *If there exists a syncretic assignment mapping every $E \in \mathcal{E}_H$ to a Horn compliant total pre-order $\leq_E$, then we can define an operator $\Delta\colon \mathcal{E}_H \times \mathcal{L}_H \rightarrow \mathcal{K}_H$ by taking $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$, for any $\mu \in \mathcal{L}_H$, and Δ satisfies postulates IC_0 – IC_8.*

4 Strengthening the postulates

Conversely, we want to show that for any Horn merging operator Δ there exists a syncretic assignment which represents it. This is true when the language is not restricted (see Proposition 1), but interesting problems arise as soon as we restrict ourselves to the Horn case.

We know from existing work on Horn revision [Delgrande and Peppas, 2015] that we can find non-syncretic assignments to represent a Horn operator Δ. Such assignments are non-syncretic in the sense that they contain non-transitive cycles between interpretations. Yet, we can still define an operator on top of these rankings, which satisfies postulates IC_0 – IC_8. Furthermore, it can be shown that there are no non-cyclic pre-orders that represent the same operator Δ. The solution proposed in [Delgrande and Peppas, 2015] is to add an extra postulate, called Acyc, specifically to eliminate cycles:

(Acyc) If for every $n \geq 1$ and $i \in \{0, n-1\}$, $\mu_i \wedge \Delta_{\mu_{i+1}}(E)$ and $\mu_n \wedge \Delta_{\mu_0}(E)$ are all consistent, then $\mu_0 \wedge \Delta_{\mu_n}(E)$ is also consistent.

Acyc provably follows from postulates IC_0 – IC_8 in full propositional logic, so it only makes a difference when we restrict the language to the Horn fragment. Here we employ the same strategy of adding an extra postulate to deal with non-transitive cycles, but we propose a postulate formulated in terms of complete formulas:

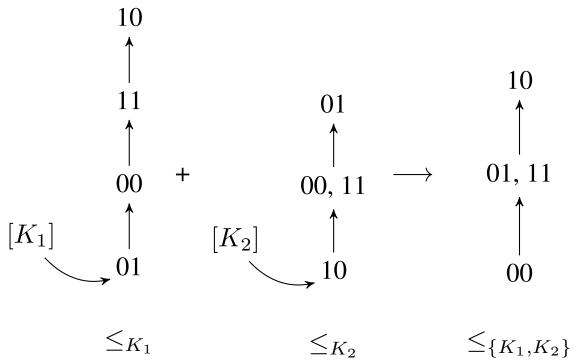

Figure 1: s_4 does not hold.

(Acyc′) For any complete formulas $\sigma_0, \dots, \sigma_n$, $n \geq 1$ and $i \in \{0, n-1\}$, it holds that if $\sigma_i \wedge \Delta_{\sigma_{i,i+1}}(E)$ and $\sigma_n \wedge \Delta_{\sigma_{n,0}}(E)$ are all consistent, then $\sigma_0 \wedge \Delta_{\sigma_{0,n}}(E)$ is also consistent.

There is clearly a strong similarity between Acyc and Acyc′, though we prefer the latter here for its intuitive appeal. Moreover, it can be shown that they are equivalent modulo the merging postulates.

Proposition 2. *Given postulates* IC_0 – IC_8*,* Acyc *and* Acyc′ *are equivalent.*

The intuition behind Acyc′ is that it prevents non-transitive cycles between chains of interpretations of arbitrary length. Suppose $n = 2$ and the antecedent of Acyc′ is true: then from the fact that $\sigma_0 \wedge \Delta_{\sigma_{0,1}}(E)$ is consistent we conclude that $w_0 \in [\Delta_{\sigma_{0,1}}(E)]$, where $[\sigma_i] = \{w_i\}$. This means that w_0 is among the models of $\sigma_{0,1}$ that are 'preferred', or considered more plausible, by Δ. Thus, in the pre-order $\leq_E$ that represents Δ, it should hold that $w_0 \leq_E w_1$. By the same token, we get that $w_1 \leq_E w_2 \leq_E w_0$ should hold. Since we want $\leq_E$ to be transitive, it should also hold that $w_0 \leq_E w_2$, and this is exactly what Acyc′ requires at this point. Thus, we need the extra postulate Acyc′(or something equivalent) to ensure that the pre-orders representing a given Horn merging operator Δ preserve intuitive properties such as (in this case) transitivity.

Introducing Acyc′ is not enough, as one can still find assignments that represent a Horn merging operator Δ without being syncretic, this time because they do not satisfy properties s_4 – s_6. The following examples make this clearer.

Example 4. Consider Horn knowledge bases $[K_1] = \{01\}$, $[K_2] = \{10\}$ and an assignment that works as in Figure 1 when restricted to K_1 and K_2. Figure 1 shows the rankings associated with K_1 and K_2 and the result of merging them into the new ranking $\leq_{\{K_1,K_2\}}$.[1] Notice that s_4 is not true: s_4 requires that $01 \approx_{\{K_1,K_2\}} 10$, whereas we have $01 <_{\{K_1,K_2\}} 10$. However, we can define a (Horn) merging operator Δ on top of this assignment in the usual way, by taking $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$, for any Horn formula μ, and Δ will satisfy postulates IC_0 – IC_8 + Acyc′. This can be verified by direct inspection; we focus here only on IC_4. The problematic interpretations are the models 01 and 10 of K_1 and K_2, respectively. Notice that there is no Horn formula that represents exactly the set $\{01, 10\}$: the best we can do is take a Horn formula μ such that $[\mu] = [\sigma_{01,10}] = \{00, 01, 10\}$. Obviously, $K_1 \models \mu$ and $K_2 \models \mu$, so we are in the range of application of IC_4. We have that $[\Delta_\mu(\{K_1, K_2\})] = \{00\}$, and $[\Delta_\mu(\{K_1, K_2\}) \wedge K_1] = [\Delta_\mu(\{K_1, K_2\}) \wedge K_2] = \emptyset$, so IC_4 is satisfied for this particular μ.

[1]It is worth noting that $\leq_{K_1}$, $\leq_{K_2}$ and $\leq_{\{K_1,K_2\}}$ are *not* generated using any familiar notion of distance—the rankings were hand-picked.

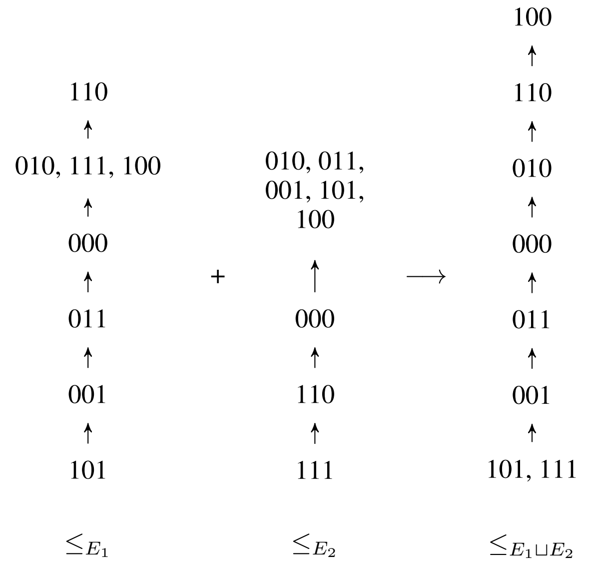

Figure 2: s_5 does not hold for 010 and 100.

Example 4 is significant because it shows that an assignment which does not satisfy s_4 may still represent an operator Δ obeying IC_4. And given the standard formulation of IC_4, this turns out to be unavoidable.

Proposition 3. *There is no syncretic assignment representing* Δ *from Example 4 that assigns to* $\{K_1, K_2\}$ *a pre-order* $\leq^\star_{\{K_1,K_2\}}$ *where* $01 \approx^\star_{\{K_1,K_2\}} 10$*.*

Proof. Suppose $01 \approx^\star_{\{K_1,K_2\}} 10$. From Figure 1 we know that $[\Delta_{\sigma_{10,11}}(\{K_1, K_2\})] = \{11\}$, hence $min_{\leq^\star}\{10, 11\} = \{11\}$ and thus $11 <^\star_{\{K_1,K_2\}} 10$. Similarly, we obtain $01 \approx^\star_{\{K_1,K_2\}} 11$, which implies $01 <^\star_{\{K_1,K_2\}} 10$. This creates a contradiction. □

The following example shows how s_5 fails to be enforced by IC_5 in the case of Horn logic.

Example 5. Assume there exists an assignment which for two profiles E_1 and E_2 behaves as in Figure 2, and is otherwise Horn compliant and syncretic. Property s_5 does not hold: $010 \approx_{E_1} 100$ and $010 \approx_{E_2} 100$, but $010 <_{E_1 \sqcup E_2} 100$. However, as in Example 4, we can define a (Horn) merging operator Δ on top of this assignment and Δ will satisfy postulates IC_0 – IC_8 + Acyc′. Let us check that IC_5 holds. The problematic interpretations here are 010 and 100 (for which s_5 does not hold). In this case we have that $\Delta_{\sigma_{010,100}}(E_1) \wedge \Delta_{\sigma_{010,100}}(E_2)$ is consistent, and

$[\Delta_{\sigma_{010,100}}(E_1) \wedge \Delta_{\sigma_{010,100}}(E_2)] = [\Delta_{\sigma_{010,100}}(E_1 \sqcup E_2)] = \{000\}$. This shows that for the case we are interested in ($\mu = \sigma_{010,100}$) IC_5 is true.

Similarly as for IC_4, we can show that such a counter-example to s_5 is unavoidable.

It is perhaps surprising to see that IC_5 can be satisfied in an assignment where s_5 does not hold, but closer thought shows this is to be expected: since in the Horn fragment we cannot represent the set $\{100, 010\}$ with a formula, it becomes harder to control the order in which 100 and 010 appear. Without any additional constraints on Δ, one cannot prevent it from varying the order of 100 and 010 in ways that directly contradict s_5. Similar counter-examples can be constructed for s_6.

Examples 4 and 5 show that a syncretic assignment with total pre-orders is not the most natural way to represent a Horn merging operator. Hence, we introduce the following notion.

Definition 3. A pre-order $\leq$ on $\mathcal{W}$ is *Horn connected* if

(h_1) $\leq$ is Horn compliant,

(h_2) any $w_i, w_j \in \mathcal{W}$ that are in the subset relation are in $\leq$, and

(h_3) for any $w_i, w_j \in \mathcal{W}$ such that $w_i \nsubseteq w_j$ and $w_j \nsubseteq w_i$, it holds that if $w_i \leq w_j$ then:

- ($\mathsf{h}_{3.1}$) $w_i \in min_{\leq} Cl_{\cap}(\{w_i, w_j\})$, or
- ($\mathsf{h}_{3.2}$) for some $n > 2$, there exist pair-wise distinct interpretations $w_1, \ldots, w_n$, such that $w_1 = w_i$, $w_n = w_j$ and $w_1 \leq \cdots \leq w_n$.

A Horn connected pre-order $\leq_E$ is not necessarily total. Example 6 illustrates this.

Example 6. Consider the following pre-orders on the 2-letter alphabet: (a) $11 <_1 01 <_1 10 <_1 00$, (b) $00 <_2 01 <_2 11 <_2 10$, (c) $00 <_3 01 <_3 11, 00 <_3 10 <_3 11$, $01 \nleq_3 10$ $10 \nleq_3 01$, (d) $00 <_4 01 \approx_4 10 <_4 11$. It is immediately visible that all pre-orders are Horn compliant (h_1) and that they satisfy h_2. Let us focus on interpretations 01 and 10.

In $\leq_1$ we have $01 \in min_{\leq} Cl_{\cap}(\{01, 10\})$. Thus $\mathsf{h}_{3.1}$ is satisfied, and $\leq_1$ is Horn connected. In $\leq_2$ we do not have $01 \in min_{\leq} Cl_{\cap}(\{01, 10\})$, but there is interpretation 11 such that $01 < 11 < 10$. Thus $\mathsf{h}_{3.2}$ is satisfied and $\leq_2$ is Horn connected. Pre-order $\leq_3$ is partial, as 01 and 10 are not in $\leq_3$, and thus h_3 is vacuously true. In $\leq_4$ we have $01 \approx 10$ though none of $\mathsf{h}_{3.1}$ and $\mathsf{h}_{3.2}$ holds, thus $\leq_4$ is not Horn connected.

Next, for every Horn operator Δ and Horn profile E, we define a (partial) pre-order on complete formulas of $\mathcal{L}_H$.[2]

Definition 4. Given a Horn operator Δ, then for any Horn profile E and complete Horn formulas σ_i, σ_j, we say that $\sigma_i \preceq_E \sigma_j$ if there exist complete Horn formulas $\sigma_1, \ldots, \sigma_n$ such that $\sigma_1 = \sigma_i$, $\sigma_n = \sigma_j$, and for $i \in \{1, n-1\}$, $\sigma_i \wedge \Delta_{\sigma_{i,i+1}}(E)$ are all consistent.

It is straightforward to check that $\preceq_E$ is reflexive and transitive, and thus a pre-order on complete Horn formulas. We write $\prec_E$ for the strict part of $\preceq_E$. It is also worth noting that $\preceq_E$ is total when the underlying language is full propositional logic, since $[\sigma_{i,j}] = \{w_i, w_j\}$ and we can take the sequence $\sigma_1, \ldots, \sigma_n$ to be just σ_i, σ_j or σ_j, σ_i. This does not necessarily hold in the case of the Horn fragment, where $\preceq_E$ can be partial.

We now reformulate IC_4, IC_5 and IC_6 for $K_1, K_2 \in \mathcal{K}_H$, $E_1, E_2 \in \mathcal{E}_H$ and complete Horn formulas σ_i, σ_j as follows:

(IC'_4) For any $\sigma_i \models K_1$, there exists $\sigma_j \models K_2$ such that $\sigma_j \preceq_{\{K_1, K_2\}} \sigma_i$.

(IC'_5) If $\sigma_i \preceq_{E_1} \sigma_j$ and $\sigma_i \preceq_{E_2} \sigma_j$, then $\sigma_i \preceq_{E_1 \sqcup E_2} \sigma_j$.

(IC'_6) If $\sigma_i \preceq_{E_1} \sigma_j$ and $\sigma_i \prec_{E_2} \sigma_j$, then $\sigma_i \prec_{E_1 \sqcup E_2} \sigma_j$.

These postulates make a difference only in the Horn fragment, while in full propositional logic they are redundant.

Proposition 4. *In the case of full propositional logic,* IC'_4, IC'_5 *and* IC'_6 *follow from the standard* $\mathsf{IC}_0 - \mathsf{IC}_8$ *postulates.*

With postulates IC'_4, IC'_5 and IC'_6 we can derive a representation result for syncretic assignments with Horn connected pre-orders. The result is split across two theorems: Theorem 3 shows that Horn connected pre-orders can be used to construct a Horn merging operator satisfying our amended set of postulates. Its converse, Theorem 4, shows that any Horn merging operator satisfying the amended postulates is represented by a syncretic assignment with Horn connected pre-orders.

Theorem 3. *If there exists a syncretic assignment mapping every* $E \in \mathcal{E}_H$ *to a Horn connected total pre-order* $\leq_E$*, then we can define an operator* $\Delta\colon \mathcal{E}_H \times \mathcal{L}_H \to \mathcal{K}_H$ *by taking* $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$*, for any* $\mu \in \mathcal{L}_H$*, and* Δ *satisfies postulates* $\mathsf{IC}_0 - \mathsf{IC}_3 + \mathsf{IC}'_4 + \mathsf{IC}'_5 + \mathsf{IC}'_6 + \mathsf{IC}_7 - \mathsf{IC}_8 + \mathsf{Acyc}'$.

Theorem 4. *If a Horn operator* $\Delta\colon \mathcal{E}_H \times \mathcal{L}_H \to \mathcal{K}_H$ *satisfies postulates* $\mathsf{IC}_0 - \mathsf{IC}_3 + \mathsf{IC}'_4 + \mathsf{IC}'_5 + \mathsf{IC}'_6 + \mathsf{IC}_7 - \mathsf{IC}_8 + \mathsf{Acyc}'$*, then there exists a syncretic assignment mapping every Horn profile* E *to a Horn connected pre-order* $\leq_E$*, such that, for any Horn formula* μ*, it holds that* $[\Delta_\mu(E)] = min_{\leq_E}[\mu]$.

In Theorem 4, the strengthened postulates IC'_4, IC'_5 and IC'_6 rule out Horn merging operators Δ represented by non-syncretic assignments such as the ones in Examples 4 and 5, and thus justify their presence. Our focus on Horn connected pre-orders, on the other hand, should not be seen as a restriction: we can translate any Horn compliant pre-order $\leq_E$ into a Horn connected one $\leq^\star_E$ such that the overall assignment (1) represents the same (Horn) merging operator and (2) remains syncretic. This can be done simply by 'uncoupling' pairs w_i and w_j which are not in the subset relation and do not satisfy either of the properties $\mathsf{h}_{3.1}$ and $\mathsf{h}_{3.2}$. Since w_i and w_j do not appear together in any set of the type $min_{\leq_E}[\mu]$, for $\mu \in \mathcal{L}_H$, the Horn merging operators represented by $\leq_E$ and $\leq^\star_E$ are the same.

However, the reverse is not as straightforward: for any Horn connected pre-order there exist more than one Horn compliant pre-orders representing the same Horn merging operator: any interpretations w_i and w_j that are not in $\leq^\star_E$ can be related in several ways if we care about making $\leq^\star_E$ total (we could have $w_i <^\star_E w_j$, or $w_j <^\star_E w_i$, etc.), and some of the configurations give rise to non-syncretic assignments. Our point is that w_i and w_j do not need to be related as long as the represented merging operator Δ stays the same. Indeed, the main motivation for

[2]The pre-order $\preceq_E$ is not to be confused with the pre-order $\leq_E$ on interpretations, though it is meant to mirror it.

formulating the representation result with partial pre-orders is that if w_i and w_j do not satisfy $\mathsf{h}_{3.1}$ and $\mathsf{h}_{3.2}$ then a Horn merging operator Δ does not give us any information on what the order between them should be. It makes sense, in this case, to not include w_i and w_j in the pre-order representing Δ.

5 A concrete Horn merging operator

By Theorem 2, we can find a Horn merging operator simply by exhibiting a Horn compliant, syncretic assignment. As in Examples 1 and 2, we can specify a pre-order $\leq_K$ by assigning numbers to interpretations, relative to K, and in the rest of this section this is how we will be thinking of pre-orders. We write $l_K(w)$ to denote the number assigned to w with respect to some knowledge base K. If K has exactly one model w', we simply write $l_{w'}(w)$.

One difficulty here is that there is no obvious candidate for an off-the-shelf assignment that satisfies all the required properties: Horn compliance rules out standard approaches using familiar distances between interpretations. Therefore, we start by describing some general conditions sufficient to guarantee that the resulting assignment satisfies $\mathsf{s}_1 - \mathsf{s}_6$ and is Horn compliant.

We take $l_K(w) \geq 0$, for any knowledge base K and any $w \in \mathcal{W}$, with $l_K(w) = 0$ if and only if $w \in [K]$. This guarantees the assignment satisfies $\mathsf{s}_1 - \mathsf{s}_3$. We use the sum Σ to aggregate individual pre-orders, and this guarantees $\mathsf{s}_5 - \mathsf{s}_6$. The next conditions spell out what is needed for an assignment to satisfy s_4.

Definition 5. The *distance between knowledge bases* K_1 *and* K_2 is defined as $d(K_1, K_2) = min\{l_{K_1}(w) \mid w \in [K_2]\}$.

We are interested in knowledge bases that satisfy the following property.

Definition 6. Knowledge bases K_1 and K_2 are *symmetric* if $d(K_1, K_2) = d(K_2, K_1)$.

Symmetry is important because it guarantees s_4.

Proposition 5 ([Konieczny and Pino Pérez, 2002]). *If an assignment satisfies* $\mathsf{s}_1 - \mathsf{s}_3$*, then it satisfies* s_4 *iff any two knowledge bases are symmetric.*

Interestingly, it turns out that if we fix the pre-orders for every knowledge base K that has exactly one model, then pre-orders for knowledge bases K with more than one model are completely determined by this initial assignment (see Example 7). Thus, if symmetry is enforced, we can represent an assignment by just giving the pre-orders for single-model knowledge bases, as a $2^n \times 2^n$ matrix. We shall call this *the initial matrix*. The same symmetry condition forces the initial matrix to be symmetric. For $\mathsf{s}_1 - \mathsf{s}_3$ to hold, the initial matrix needs to have positive entries and 0 on the main diagonal.[3]

Example 7. Table 3 shows the initial matrix for the 2 letter alphabet, plus an additional ranking obtained through symmetry. Each column represents a ranking: for instance the first column represents the ranking for a knowledge base that has 00 as its sole model. The number assigned to 00 in this ranking is 0, the number assigned to 01 is 1, etc. The ranking for a knowledge base K that has $\{10, 11\}$ as its set of models is computed from the initial assignment matrix with symmetry. For example, consider interpretation 00. By symmetry, we have that $l_K(00) = l_{00}(K)$. Thus, we obtain $l_{00}(K) = min\{l_{00}(10), l_{00}(11)\} = min\{2, 3\} = 2$.

[3] All this is treated rigorously in [Haret *et al.*, 2015].

	00	01	10	11	$\{10, 11\}$
00	0	1	2	3	2
01	1	0	3	5	3
10	2	3	0	8	0
11	3	5	8	0	0

Table 3: An initial assignment determines the remaining rankings by symmetry.

All that is left is Horn compliance, and here we propose the following notion.

Definition 7. A pre-order $\leq$ is *well-behaved* if and only if for any interpretations w_0, w_1, w_2 such that $w_1 \not\subseteq w_2$, $w_2 \not\subseteq w_1$ and $w_0 = w_1 \cap w_2$, it is the case that $w_0 \leq w_1$ or $w_0 \leq w_2$ and $|min\{l(w_1), l(w_2)\} - l(w_0)| \leq |max\{l(w_1), l(w_2)\} - l(w_0)|$.

Notice that a well-behaved pre-order $\leq$ is also Horn compliant. What makes well-behavedness suitable for our needs, however, is that it is transmitted through Σ-aggregation.

Proposition 6. *If* $\leq_1$ *and* $\leq_2$ *are well-behaved, then the pre-order obtained by* Σ*-aggregating* $\leq_1$ *and* $\leq_2$ *is well-behaved.*

Using all this knowledge, we can now define a specific Horn compliant syncretic assignment, which we will call *the summation assignment*. We define this assignment for the general case of an alphabet of size n. As suggested by the previous discussion, we give the initial matrix and use symmetry to determine pre-orders $\leq_K$, when $|[K]| > 1$. Since the matrix for the initial assignment has to itself be symmetric and have 0 on the main diagonal, we will only define the entries in the matrix below the main diagonal, with the understanding that the entries above the main diagonal are fixed by symmetry. Also, the order in which interpretations appear in the rows and columns is fixed by the number of 1's in the corresponding bit-vector. For instance, the matrix for the 3-letter alphabet has its rows and columns ordered as follows: 000, 001, 010, 100, 011, 101, 110, 111.

The definition of the bottom half of the initial assignment matrix is recursive. First, put:

$$l_{w_0}(w_i) = i, \text{ for } i \in \{0, \ldots, 2^n - 1\}.$$

Hence, the levels on the first column are $0, 1, 2, \ldots, 2^n - 1$ (see Table 4).

Second, for $1 \leq i \leq 2^n - 1$, put:

$$l_{w_i}(w_{i+1}) = l_{w_{i-1}}(w_i) + l_{w_{i-1}}(w_{i+1}).$$

Roughly, this means that the number in a particular cell under the main diagonal is the sum of its two neighbours to the left. In Table 4: if $l_{w_{i-1}}(w_i) = a$, $l_{w_{i-1}}(w_{i+1}) = b$, then $l_{w_i}(w_{i+1}) = a + b$. This is simpler than it sounds, and Table 3 shows the matrix that we get for the 2-letter alphabet.

The following proposition guarantees that the summation assignment is Horn compliant and stays Horn compliant through repeated Σ-aggregations.

	w_0	$\dots$	w_{i-1}	w_i	w_{i+1}	$\dots$
w_0	0	$\dots$	$i-1$	i	$i+1$	$\dots$
$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$
w_{i-1}	$i-1$	$\dots$	0			$\dots$
w_i	i	$\dots$	a	0		$\dots$
w_{i+1}	$i+1$	$\dots$	b	$a+b$	0	$\dots$
$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$	$\dots$

Table 4: The recursive relation for levels.

Proposition 7. *The summation assignment is well-behaved.*

This is the last piece of information needed. We can now assert the following theorem.

Theorem 5. *The summation assignment represents a Horn merging operator.*

6 Conclusion and future work

In this paper, we provided a novel representation theorem for Horn merging by strengthening the standard merging postulates. Belief change operators for the Horn fragment have attracted increasing attention over the last years, in particular revision and contraction, while merging in the Horn fragment remained rather unexplored so far. An exception is the work by Creignou *et al.* [2014], who proposed to *adapt* known merging operators by means of a certain post-processing and studied the limits of this approach in terms of satisfaction of the merging postulates. One of the main results of that paper is that in their framework it is not possible to keep all postulates satisfied. In our work, we have presented a novel concrete Horn merging operator satisfying *all* postulates.

The moral of the present work is that, while going from syncretic assignments to Horn merging operators is relatively easy (Horn compliance is sufficient, by Theorem 2), going from Horn merging operators to syncretic assignments requires considerably more machinery (in particular, stronger postulates). Thus, all the work in Section 4 is needed to obtain a full representation result. Even so, Section 5 highlights that the easiness of the first direction is only relative, as finding concrete syncretic assignments that are also Horn compliant requires some conceptual work, and there is no obvious trivial operator that does the job. The main difficulty here lies in making sure that if two pre-orders $\leq_1$ and $\leq_2$ are Horn compliant, then the pre-order resulted from Σ-aggregating them is also Horn compliant. Our well-behavedness property guarantees this.

Future work on merging in the Horn fragment would have to consider extending the family of Horn merging operators. This requires seeing how Horn compliance interacts, on the model side, with other aggregation functions (such as $GMAX$) and exploring the range of conditions guaranteeing Horn compliance of an assignment. We may add to this the study of other merging postulates (e.g., majority and arbitration), considered in the merging literature [Konieczny and Pino Pérez, 2002; 2011] but not touched upon here. Finally, we would like to extend our approach to other fragments of propositional logic (e.g., Krom or dual Horn), where similar problems arise and for which tailored notions of compliance and strengthened postulates are likely needed.

References

[Alchourrón *et al.*, 1985] 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.

[Booth *et al.*, 2011] Richard Booth, Thomas A. Meyer, Ivan J. Varzinczak, and Renata Wassermann. On the link between partial meet, kernel, and infra contraction and its application to Horn logic. *J. Artif. Intell. Res. (JAIR)*, 42:31–53, 2011.

[Creignou *et al.*, 2014] Nadia Creignou, Odile Papini, Stefan Rümmele, and Stefan Woltran. Belief merging within fragments of propositional logic. In *Proc. ECAI 2014*, volume 263 of *Frontiers in Artificial Intelligence and Applications*, pages 231–236. IOS Press, 2014.

[Delgrande and Peppas, 2015] James P. Delgrande and Pavlos Peppas. Belief revision in Horn theories. *Artif. Intell.*, 218:1–22, 2015.

[Delgrande and Wassermann, 2013] James P. Delgrande and Renata Wassermann. Horn clause contraction functions. *J. Artif. Intell. Res. (JAIR)*, 48:475–511, 2013.

[Haret *et al.*, 2015] Adrian Haret, Stefan Rümmele, and Stefan Woltran. Merging in the Horn Fragment. Technical Report DBAI-TR-2015-91, Technische Universität Wien, 2015.

[Katsuno and Mendelzon, 1991] Hirofumi Katsuno and Alberto O. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proc. KR 1991*, pages 387–394. Morgan Kaufmann, 1991.

[Katsuno and Mendelzon, 1992] Hirofumi Katsuno and Alberto O. Mendelzon. Propositional knowledge base revision and minimal change. *Artif. Intell.*, 52(3):263–294, 1992.

[Konieczny and Pino Pérez, 2002] Sébastien Konieczny and Ramón Pino Pérez. Merging information under constraints: A logical framework. *J. Log. Comput.*, 12(5):773–808, 2002.

[Konieczny and Pino Pérez, 2011] Sébastien Konieczny and Ramón Pino Pérez. Logic based merging. *J. Philosophical Logic*, 40(2):239–270, 2011.

[Schaefer, 1978] Thomas J. Schaefer. The complexity of satisfiability problems. In *Proc. STOC 1978*, pages 216–226. ACM, 1978.

[Van De Putte, 2013] Frederik Van De Putte. Prime implicates and relevant belief revision. *J. Log. Comput.*, 23(1):109–119, 2013.

[Zhuang and Pagnucco, 2012] Zhi Qiang Zhuang and Maurice Pagnucco. Model based horn contraction. In *Proc. KR 2012*. AAAI Press, 2012.

[Zhuang *et al.*, 2013] Zhi Qiang Zhuang, Maurice Pagnucco, and Yan Zhang. Definability of Horn revision from Horn contraction. In *Proc. IJCAI 2013*, 2013.