

Dynamics of Beliefs

Sébastien Konieczny

CRIL - CNRS
Université d'Artois, Lens, France
konieczny@cril.fr

Abstract. The dynamics of beliefs is one of the major components of any autonomous system, that should be able to incorporate new pieces of information. In this paper we give a quick overview of the main operators for belief change, in particular revision, update, and merging, when the beliefs are represented in propositional logic. And we discuss some works on belief change in more expressive frameworks.

1 Introduction

Every autonomous agent has to use a belief base to model the state of the world. This information is precious since beliefs can be costly to obtain and since they are necessary to carry on reasoning tasks or to take the appropriate decisions. So a first class requirement in order to design intelligent autonomous agents is to try to provide her the means to obtain, and to maintain the most faithful belief base. In particular an agent has to be able to incorporate new pieces of information, and to correct the incorrect beliefs when she detected them. So this dynamics of beliefs is one of the major components of any autonomous agent.

The aim of this paper is to recall the definition of the main belief change operators and the links between them. We focus on the classical case, where the beliefs of the agents are represented using propositional logic, before discussing some extensions in other representational frameworks.

This is a very quick presentation of belief change theory. For a complete introduction the reader should refer to the seminal books on belief revision [35, 36, 41, 64], or the recent special issue of Journal of Philosophical Logic on the 25 Years of AGM Theory [34].

2 Preliminaries

We consider a propositional language \mathcal{L} defined from a finite set of propositional variables \mathcal{P} and the standard connectives.

An interpretation ω is a total function from \mathcal{P} to $\{0, 1\}$. The set of all interpretations is denoted by \mathcal{W} . An interpretation ω is a model of a formula $\varphi \in \mathcal{L}$ if and only if it makes it true in the usual truth functional way. $mod(\varphi)$ denotes the set of models of the formula φ , i.e., $mod(\varphi) = \{\omega \in \mathcal{W} \mid \omega \models \varphi\}$. When M is a set of models we denote by φ_M a formula such that $mod(\varphi_M) = M$.

A *belief base* K is a finite set of propositional formulae. In order to simplify the notations we identify the base K with the formula φ which is the conjunction of the formulae of K ¹.

3 Revision

Belief revision aims at changing the status of some beliefs in the base that are contradicted by a more reliable piece of information. Several principles are governing this revision operation:

- First is the primacy of update principle: the new piece of information has to be accepted in the belief base after the revision. This is due to the hypothesis that the new piece of information is more reliable than the current beliefs of the agent².
- Second is the principle of coherence: the new belief base after the revision should be a consistent belief base. Asking the beliefs of the agent to be consistent is a natural requirement if one wants the agent to conduct reasoning tasks from her belief base.
- Third is the principle of minimal change: the new belief base after the revision should be as close as possible from the current belief base of the agent. This important principle aims at ensuring that no unnecessary information (noise) is added to the beliefs of the agent during the revision process, and that no unnecessary information is lost during the process: information/beliefs are usually costly to obtain, we do not want to throw them away without any serious reason.

Alchourrón, Gärdenfors and Makinson [2] proposed some postulates in order to formalize these principles for belief revision.

Definition 1 ([48]). *Let φ and μ be two formulas denoting respectively the belief base of the agent, and a new piece of information. Then $\varphi \circ \mu$ is a formula representing the new belief base of the agent. An operator \circ is an AGM belief revision operator if it satisfies the following properties:*

- (R1) $\varphi \circ \mu \vdash \mu$
- (R2) *If $\varphi \wedge \mu$ is consistent then $\varphi \circ \mu \equiv \varphi \wedge \mu$*
- (R3) *If μ is consistent then $\varphi \circ \mu$ is consistent*
- (R4) *If $\varphi_1 \equiv \varphi_2$ and $\mu_1 \equiv \mu_2$ then $\varphi_1 \circ \mu_1 \equiv \varphi_2 \circ \mu_2$*
- (R5) $(\varphi \circ \mu) \wedge \phi \vdash \varphi \circ (\mu \wedge \phi)$
- (R6) *If $(\varphi \circ \mu) \wedge \phi$ is consistent then $\varphi \circ (\mu \wedge \phi) \vdash (\varphi \circ \mu) \wedge \phi$*

When one works with a finite propositional language the above postulates, proposed by Katsuno and Mendelzon, are equivalent to AGM ones [2, 35].

(R1) states that the new piece of information must be believed after the revision. (R2) says that when there is no conflict between the new piece of information and the

¹ Some approaches are sensitive to syntactical representation. In that case it is important to distinguish between K and the conjunction of its formulae (see e.g. [52]).

² If this is not the case one should use a non-prioritized revision operator [42] or a merging operator (see Section 5).

current beliefs of the agent, the revision is just the conjunction. (R3) says that revision always lead to a consistent belief base, unless the new piece of information is not consistent. (R4) is an irrelevance of syntax condition, it states that logically equivalent bases must lead to the same result. (R5) and (R6) give conditions on the revision by a conjunction.

Alchourrón, Gärdenfors and Makinson also defined contraction operators, that aim to remove some piece of information from the beliefs of the agent. These contraction operators are closely related to revision operators, since each contraction operator can be used to define a revision operator, through the Levy identity and conversely each revision operator can be used to define a contraction operator through the Harper identity [2, 35]. So one can study indifferently revision or contraction operators. So we focus on revision here.

Several representation theorems, that give constructive ways to define AGM revision/contraction operators, have been proposed, such as partial meet contraction/revision [2], epistemic entrenchments [37, 35], safe contraction [1], etc. In [48], Katsuno and Mendelzon give a representation theorem, showing that each revision operator corresponds to a faithful assignment, that associates to each base a plausibility preorder on interpretations (this idea can be traced back to Grove systems of spheres [40]).

Definition 2. *A faithful assignment is a function mapping each base φ to a pre-order \leq_φ over interpretations such that:*

1. *If $\omega \models \varphi$ and $\omega' \models \varphi$, then $\omega \simeq_\varphi \omega'$*
2. *If $\omega \models \varphi$ and $\omega' \not\models \varphi$, then $\omega <_\varphi \omega'$*
3. *If $\varphi \equiv \varphi'$, then $\leq_\varphi = \leq_{\varphi'}$*

Theorem 1 ([48]). *An operator \circ is an AGM revision operator (ie. it satisfies (R1)-(R6)) if and only if there exists a faithful assignment that maps each base φ to a total pre-order \leq_φ such that $\text{mod}(\varphi \circ \mu) = \min(\text{mod}(\mu), \leq_\varphi)$.*

One of the main problems of this characterization of belief revision is that it does not constrain the operators enough for ensuring a good behavior when we do iteratively several revisions. So one needs to add more postulates and to represent the beliefs of the agent with a more complex structure than a simple belief base. In [26] Darwiche and Pearl proposed a convincing extension of AGM revision. This proposal have been improved by an additional condition in [17, 45]. And [55, 51] define improvement operators that are a generalization of iterated revision operators.

4 Update

Whereas belief revision should be used to improve the beliefs of the agent by incorporating more reliable pieces of evidence, belief update operators aim at maintaining the belief base of the agent up-to-date, by allowing to modify the base according to a reported change in the world. This distinction between revision and update was made clear in [47, 49], where Katsuno and Mendelzon proposed postulates for belief update.

Definition 3 ([47, 49]). An operator \diamond is a (partial) update operator if it satisfies the properties (U1)-(U8). It is a total update operator if it satisfies the properties (U1)-(U5), (U8), (U9).

- (U1) $\varphi \diamond \mu \vdash \mu$
- (U2) If $\varphi \vdash \mu$, then $\varphi \diamond \mu \equiv \varphi$
- (U3) If $\varphi \not\vdash \perp$ and $\mu \not\vdash \perp$ then $\varphi \diamond \mu \not\vdash \perp$
- (U4) If $\varphi_1 \equiv \varphi_2$ and $\mu_1 \equiv \mu_2$ then $\varphi_1 \diamond \mu_1 \equiv \varphi_2 \diamond \mu_2$
- (U5) $(\varphi \diamond \mu) \wedge \phi \vdash \varphi \diamond (\mu \wedge \phi)$
- (U6) If $\varphi \diamond \mu_1 \vdash \mu_2$ and $\varphi \diamond \mu_2 \vdash \mu_1$, then $\varphi \diamond \mu_1 \equiv \varphi \diamond \mu_2$
- (U7) If φ is a complete formula, then $(\varphi \diamond \mu_1) \wedge (\varphi \diamond \mu_2) \vdash \varphi \diamond (\mu_1 \vee \mu_2)$
- (U8) $(\varphi_1 \vee \varphi_2) \diamond \mu \equiv (\varphi_1 \diamond \mu) \vee (\varphi_2 \diamond \mu)$
- (U9) If φ is a complete formula and $(\varphi \diamond \mu) \wedge \phi \not\vdash \perp$, then $\varphi \diamond (\mu \wedge \phi) \vdash (\varphi \diamond \mu) \wedge \phi$

Most of these postulates are close to the ones of revision. The main differences lie in postulate (U2) that is much weaker than (R2): conversely to revision, even if the new piece of information is consistent with the belief base, the result is generally not simply the conjunction. This illustrates the fact that revision can be seen as a selection process of the most plausible worlds of the current beliefs with respect to the new piece of information, whereas update is a transition process: each world of the current beliefs have to be translated to the closest world allowed by the new piece of information. This world-by-world treatment is expressed by postulate (U8).

As for revision, there is a representation theorem in terms of faithful assignment.

Definition 4. A faithful assignment is a function mapping each interpretation ω to a pre-order \leq_ω over interpretations such that if $\omega \neq \omega'$, then $\omega <_\omega \omega'$.

One can easily check that this faithful assignment on interpretations is just a special case of the faithful assignment on bases defined in the previous section on the complete base corresponding to the interpretation.

Katsuno and Mendelzon give two representation theorems for update operators. The first representation theorem, that is the most commonly used, corresponds to partial pre-orders. This use of partial pre-order is one of the differences between belief revision and belief update (note nonetheless that postulates for belief revision can also be adapted to modelize assignments giving partial pre-orders [9]).

Theorem 2 ([47, 49]). An update operator \diamond satisfies (U1)-(U8) if and only if there exists a faithful assignment that maps each interpretation ω to a partial pre-order \leq_ω such that $\text{mod}(\varphi \diamond \mu) = \bigcup_{\omega \models \varphi} \min(\text{mod}(\mu), \leq_\omega)$.

But there is also a second theorem corresponding to total pre-orders.

Theorem 3 ([47, 49]). An update operator \diamond satisfies (U1)-(U5), (U8) and (U9) if and only if there exists a faithful assignment that maps each interpretation ω to a total pre-order \leq_ω such that $\text{mod}(\varphi \diamond \mu) = \bigcup_{\omega \models \varphi} \min(\text{mod}(\mu), \leq_\omega)$.

This characterization of update is quite convincing, but some criticisms can be made that suggest that more elaborate update operators can be studied [43].

5 Merging

Merging operators [4, 5, 62, 58, 56] should be used when one wants to combine several belief bases, or wants to take into account several pieces of information of same reliability.

We first need to define a profile of bases, that will represent the set of bases/information one wants to combine:

A *profile* Ψ is a non-empty multi-set (bag) of bases $\Psi = \{\varphi_1, \dots, \varphi_n\}$ (hence different agents are allowed to exhibit identical bases), and represents a group of n agents. We denote by $\bigwedge \Psi$ the conjunction of bases of $\Psi = \{\varphi_1, \dots, \varphi_n\}$, i.e., $\bigwedge \Psi = \varphi_1 \wedge \dots \wedge \varphi_n$. A profile Ψ is said to be consistent if and only if $\bigwedge \Psi$ is consistent. The multi-set union is denoted by \sqcup .

Belief merging operators aim at aggregating several bases into a single one. The most basic case is when all the bases have the same strength/importance (see [28] for a discussion on prioritized merging). Often the aggregation has to obey a set of rules, that can be a translation of physical laws or of some knowledge about the result, that form the integrity constraints for the merging. Let us see the postulates for Integrity Constraints merging operators:

Definition 5 ([53]). *Let Ψ be a profile and μ be a formula encoding integrity constraints. Then $\Delta_\mu(\Psi)$ represents the merging of the profile Ψ under the integrity constraints μ . An operator Δ is an IC merging operator if it satisfies the following properties:*

- (IC0) $\Delta_\mu(\Psi) \vdash \mu$
- (IC1) *If μ is consistent, then $\Delta_\mu(\Psi)$ is consistent*
- (IC2) *If $\bigwedge \Psi$ is consistent with μ , then $\Delta_\mu(\Psi) \equiv \bigwedge \Psi \wedge \mu$*
- (IC3) *If $\Psi_1 \equiv \Psi_2$ and $\mu_1 \equiv \mu_2$, then $\Delta_{\mu_1}(\Psi_1) \equiv \Delta_{\mu_2}(\Psi_2)$*
- (IC4) *If $\varphi_1 \vdash \mu$ and $\varphi_2 \vdash \mu$, then $\Delta_\mu(\{\varphi_1, \varphi_2\}) \wedge \varphi_1$ is consistent if and only if $\Delta_\mu(\{\varphi_1, \varphi_2\}) \wedge \varphi_2$ is consistent*
- (IC5) $\Delta_\mu(\Psi_1) \wedge \Delta_\mu(\Psi_2) \vdash \Delta_\mu(\Psi_1 \sqcup \Psi_2)$
- (IC6) *If $\Delta_\mu(\Psi_1) \wedge \Delta_\mu(\Psi_2)$ is consistent, then $\Delta_\mu(\Psi_1 \sqcup \Psi_2) \vdash \Delta_\mu(\Psi_1) \wedge \Delta_\mu(\Psi_2)$*
- (IC7) $\Delta_{\mu_1}(\Psi) \wedge \mu_2 \vdash \Delta_{\mu_1 \wedge \mu_2}(\Psi)$
- (IC8) *If $\Delta_{\mu_1}(\Psi) \wedge \mu_2$ is consistent, then $\Delta_{\mu_1 \wedge \mu_2}(\Psi) \vdash \Delta_{\mu_1}(\Psi)$*

These postulates are quite close to the ones of revision. The ones that specifically talk about aggregation are (IC4), (IC5) and (IC6). (IC4) is a fairness postulate, that expresses the fact that all the bases have the same importance/weight, so when merging two such bases one can not give more importance to one of them. (IC5) and (IC6) talk about the result of the merging when we join two groups. (IC5) states that all that is common in the merging of the two groups must be selected if we join the two groups. And (IC6) strengthen this condition by asking that the merging obtained when we join the two groups have to be exactly what is commonly chosen by the two groups. These two postulates correspond to well known Pareto conditions (see conditions 5 and 6 of the syncretic assignment).

There is also a representation theorem for merging operators in terms of pre-orders on interpretations [53].

Definition 6. A syncretic assignment is a function mapping each profile Ψ to a total pre-order \leq_{Ψ} over interpretations such that:

1. If $\omega \models \Psi$ and $\omega' \models \Psi$, then $\omega \simeq_{\Psi} \omega'$
2. If $\omega \models \Psi$ and $\omega' \not\models \Psi$, then $\omega <_{\Psi} \omega'$
3. If $\Psi_1 \equiv \Psi_2$, then $\leq_{\Psi_1} = \leq_{\Psi_2}$
4. $\forall \omega \models \varphi \exists \omega' \models \varphi' \omega' \leq_{\{\varphi\} \sqcup \{\varphi'\}} \omega$
5. If $\omega \leq_{\Psi_1} \omega'$ and $\omega \leq_{\Psi_2} \omega'$, then $\omega \leq_{\Psi_1 \sqcup \Psi_2} \omega'$
6. If $\omega <_{\Psi_1} \omega'$ and $\omega \leq_{\Psi_2} \omega'$, then $\omega <_{\Psi_1 \sqcup \Psi_2} \omega'$

Theorem 4 ([53]). An operator Δ is an IC merging operator if and only if there exists a syncretic assignment that maps each profile Ψ to a total pre-order \leq_{Ψ} such that

$$\text{mod}(\Delta_{\mu}(\Psi)) = \min(\text{mod}(\mu), \leq_{\Psi})$$

6 On the links between revision, update and merging

6.1 Revision vs Update

Intuitively revision operators bring a minimal change to the base by selecting the most plausible models among the models of the new information. Whereas update operators bring a minimal change to each possible world (model) of the base in order to take into account the change described by the new information whatever the possible world. So, if we look closely to the representation theorems (theorems 1, 2 and 3), we easily find the following result:

Theorem 5. If \circ is a revision operator (i.e. it satisfies (R1)-(R6)), then the operator \diamond defined by $\varphi \diamond \mu = \bigvee_{\omega \models \varphi} \varphi_{\{\omega\}} \circ \mu$ is an update operator that satisfies (U1)-(U9).

So this proposition states that update can be viewed as a kind of pointwise revision.

6.2 Revision vs Merging

Intuitively revision operators select in a formula (the new evidence) the closest information to a ground information (the old base). And, identically, IC merging operators select in a formula (the integrity constraints) the closest information to a ground information (a profile of bases). So following this idea it is easy to make a correspondence between IC merging operators and belief revision operators:

Theorem 6 ([53]). If Δ is an IC merging operator (it satisfies (IC0-IC8)), then the operator \circ , defined as $\varphi \circ \mu = \Delta_{\mu}(\varphi)$, is an AGM revision operator (it satisfies (R1-R6)).

See [53] for more links between belief revision and merging.

7 Other belief change operators

7.1 Confluence operators

As explained in the previous section, there are close connections between revision, update and merging. Update can be considered as a pointwise revision, and merging as a generalization of revision. So, as illustrated in Figure 1, one can define confluence operators [54] that can be considered as a pointwise merging, and as a generalization of update.



Fig. 1. Revision - Update - Merging - Confluence

Let us first define p-consistency for profiles:

Definition 7. A profile $\Psi = \{\varphi_1, \dots, \varphi_n\}$ is p-consistent if all its bases are consistent, i.e. $\forall \varphi_i \in \Psi$, φ_i is consistent.

Note that p-consistency is much weaker than consistency, the former just asks that all the bases of the profile are consistent, while the later asks that the conjunction of all the bases is consistent.

Definition 8. An operator \diamond is a confluence operator if it satisfies the following properties:

- (UC0) $\diamond_\mu(\Psi) \vdash \mu$
- (UC1) If μ is consistent and Ψ is p-consistent, then $\diamond_\mu(\Psi)$ is consistent
- (UC2) If Ψ is complete, Ψ is consistent and $\bigwedge \Psi \vdash \mu$, then $\diamond_\mu(\Psi) \equiv \bigwedge \Psi$
- (UC3) If $\Psi_1 \equiv \Psi_2$ and $\mu_1 \equiv \mu_2$, then $\diamond_{\mu_1}(\Psi_1) \equiv \diamond_{\mu_2}(\Psi_2)$
- (UC4) If φ_1 and φ_2 are complete formulae and $\varphi_1 \vdash \mu$, $\varphi_2 \vdash \mu$, then $\diamond_\mu(\{\varphi_1, \varphi_2\}) \wedge \mu$ is consistent if and only if $\diamond_\mu(\{\varphi_1, \varphi_2\}) \wedge \varphi_2$ is consistent
- (UC5) $\diamond_\mu(\Psi_1) \wedge \diamond_\mu(\Psi_2) \vdash \diamond_\mu(\Psi_1 \sqcup \Psi_2)$
- (UC6) If Ψ_1 and Ψ_2 are complete profiles and $\diamond_\mu(\Psi_1) \wedge \diamond_\mu(\Psi_2)$ is consistent, then $\diamond_\mu(\Psi_1 \sqcup \Psi_2) \vdash \diamond_\mu(\Psi_1) \wedge \diamond_\mu(\Psi_2)$
- (UC7) $\diamond_{\mu_1}(\Psi) \wedge \mu_2 \vdash \diamond_{\mu_1 \wedge \mu_2}(\Psi)$
- (UC8) If Ψ is a complete profile and if $\diamond_{\mu_1}(\Psi) \wedge \mu_2$ is consistent, then $\diamond_{\mu_1 \wedge \mu_2}(\Psi) \vdash \diamond_{\mu_1}(\Psi) \wedge \mu_2$
- (UC9) $\diamond_\mu(\Psi \sqcup \{\varphi \vee \varphi'\}) \equiv \diamond_\mu(\Psi \sqcup \{\varphi\}) \vee \diamond_\mu(\Psi \sqcup \{\varphi'\})$

See [54] for a representation theorem in terms of assignment for confluence operators. We just give the two results that show how confluence relates with respect to merging and update [54]:

Theorem 7. If \diamond is a confluence operator (i.e. it satisfies (UC0-UC9)), then the operator \diamond , defined as $\varphi \diamond \mu = \diamond_\mu(\varphi)$, is a total update operator (i.e. it satisfies (U1-U9)).

For relating confluence and merging, we need to use the notion of state:

Definition 9. A multi-set of interpretations will be called a state. We use the letter e , possibly with subscripts, for denoting states. If $\Psi = \{\varphi_1, \dots, \varphi_n\}$ is a profile and $e = \{\omega_1, \dots, \omega_n\}$ is a state such that $\forall i \omega_i \models \varphi_i$, we say that e is a state of the profile Ψ , or that the state e models the profile Ψ , that will be denoted by $e \models \Psi$.

If $e = \{\omega_1, \dots, \omega_n\}$ is a state, we define the profile Ψ_e by putting $\Psi_e = \{\varphi_{\{\omega_1\}}, \dots, \varphi_{\{\omega_n\}}\}$.

Theorem 8. If Δ is an IC merging operator (i.e. it satisfies (IC0-IC8)) then the operator \diamond defined by $\diamond_\mu(\Psi) = \bigvee_{e \models \Psi} \Delta_\mu(\Psi_e)$ is a confluence operator (i.e. it satisfies (UC0-UC9)).

7.2 Extrapolation and approaches based on sequences of observations

In [31, 32] Dupin and Lang defined extrapolation operators. The idea is, from a sequence of observations at different time points, to try to find the scenarios that best explain the sequence. The principle of minimal change is translated in an inertial assumption, that states that the value of a propositional variable does not change if no change occur. We do not have direct information about the changes, but the observations at different time points inform us on such changes. So, very roughly, these operators can be seen as looking for the most plausible histories compatible with a sequence of observations and minimal change assumptions.

There are others works that deal with sequences of observations such as [57, 44] for instance. An interesting operator was proposed by Booth and Nitka [20]. It can be seen as a third-party counterpart of extrapolation. The idea is that an observer observe a sequence of inputs that receives a given agent and a sequence of corresponding outputs (parts of the belief of the agent at that time point). Then the problem is to try to identify the initial beliefs of the agent and her beliefs during the sequence.

7.3 Belief negotiation

In [16] Booth proposes to aggregate the beliefs of different agents by using a iterative selection-weakening process. The idea is, until the conjunction of the bases is consistent, to select some bases that have to weaken their beliefs. Like belief merging, these belief negotiation operators allow to obtain a consistent belief base from a set of jointly inconsistent bases. But the aim is quite different. In belief merging the aim is to extract as much information as possible from the set of bases, whereas in belief negotiation the aim is to find a potential consensual issue in a (abstract) negotiation process. Several works have tried to use tools from belief change theory in order to modelize abstract negotiation processes [14–16, 68, 59, 50, 38]. We think that there is still a lot to do in this direction. In particular there is no representation theorem for abstract negotiation.

7.4 Prioritized merging operators

In [28] Delgrande, Dubois and Lang propose an interesting discussion on prioritized merging operators. The idea is to merge a set of weighted formulae. The weights are

used to stratify the formulae (a formula with a greater weight is more important, even if they are a large number of formula with smaller weights that contradict it).

Delgrande, Dubois and Lang motivate the generality of their approach by showing that classical merging operators (on unweighted formulae) and iterated belief revision operators (à la Darwiche and Pearl [26]) can be considered as two extreme cases of this weighted merging framework.

The main argument is that if one makes the hypothesis that the new pieces of information that come successively in an iterated revision process are about a static world (the usual hypothesis), then there is no reason to give the preference to the last ones. If these information have different reliability, then this can be represented explicitly with the weights of the formulae, in order to take this difference of reliability in the iterated “revision” process if they do not come in the order corresponding to their relative reliability. And the correct way to do that is to make a prioritized merging.

This discussion is interesting since in several papers on iterated revisions, it seems that the authors do not make any distinction between the hypothesis to have more and more recent pieces of information, and the hypothesis to have more and more reliable pieces of information.

The framework of Delgrande, Dubois and Lang identifies the epistemic states as the sequences of formulae that the agent receives. They show that the postulates for iterated belief revision can be obtain as special case of their postulates for weighted merging, and that they can also lead to some postulates of IC merging. This work is interesting since it opens a way for logical characterization of prioritized merging. It could be interesting to try to find a representation theorem in this case, and to look at the generalization of IC merging operators in this prioritized merging framework.

8 Belief change in other representational frameworks

8.1 Dynamics of Horn bases

Recently some works have focus on the contraction of Horn bases [27, 18, 19, 29]. This is an interesting case since Horn bases are used for instance for deductive databases and logic programming. Usually works on belief change suppose that the logic is at least as strong as classical propositional logic. But these works on Horn bases show that restrictions of propositional logic exhibit some interesting characteristics. In particular constructions that lead to equivalent classes of operators in the classical case, give rise to different ones for the Horn case.

8.2 Merging of first order bases

Lang and Bloch propose to define model-based merging operators using the maximum as aggregation function ($\Delta^{d,\max}$) by using dilation³ process [12]. One can note that in the original Dalal paper [25], he defines his revision operator with such a dilation function rather than with a distance.

³ Roughly speaking dilation allows to reach the points/worlds in the neighborhood of a point/world. See [12] to see how to define this formally.

Gorogiannis and Hunter [39] extend this approach in order to define others model-based merging operators using dilations. So, in addition to $\Delta^{d, \max}$, they define $\Delta^{d, \Sigma}$, $\Delta^{d, \text{GMAX}}$ and $\Delta^{d, \text{GMIN}}$ operators.

The interest of this definition of these operators is that it can be easily extended to first order logic. The usual definition of model-based merging operators is based on the computation of distances between interpretations. So when using logics where the number of interpretations is infinite, this approach is not the more appropriate. The interest of defining these operators with dilations is that they can also be used in this case. This only needs to use the good dilation function. See [39] for a discussion and some examples of dilation functions in the first order logic case.

8.3 Merging of qualitative constraint networks

Condotta, Kaci, Marquis and Schwind studied the merging of qualitative constraint networks [22, 21]. These methods can be useful for merging constraint networks that represent spatial regions, for instance for Geographical Information Systems it can be necessary to merge spatial databases that come from different sources.

Conflicts that arise in this framework are more subtle than the binary ones in the propositional framework. In this case conflicts can be more or less important. For instance, if we use the Allen algebra, that allows to represent spatial information on segments on a line, namely relations as A BEFORE B, A AFTER⁴ B, A MEET B among others. A conflict between sentences A BEFORE B and A MEET B seems much less important than the one between A BEFORE B and A AFTER B.

This “intensity” that we feel between conflicts allows to define more various merging policies than in the propositional framework.

One can also look at [61, 23] to see two examples of merging of spatial regions using logical representations.

8.4 Dynamics of argumentation frameworks

There are a lot of works on argumentation as a way to reason about contradictory pieces of information. The basic idea is to use a set of arguments and an attack relation between relations. This is the starting point of Dung abstract argumentation framework [30]. In [24] the problem of merging of argumentation frameworks, where the arguments are distributed among several agents, have been studied. This requires to define a new representation frameworks for argumentation: Partial Argumentation Frameworks, where there are three possible relations between two arguments A and B. Either the agent believes that A attacks B, or he believes that A does not attacks B, or he does not know if A attacks B or not. This last case is necessary to represent the fact that an agent ignores a given argument.

The problem of revision of argumentation systems as been addressed also in several works, such as [33, 63, 13] for instance.

We think that for both argumentation revision and merging a lot of work is still necessary in order to reach convincing models.

⁴ i.e. B BEFORE A

9 Conclusion

We proposed a quick tour of the theory of belief change in classical propositional logic. The core of this theory is quite established now, with a set of important belief change operators that are logically characterized. Still, a lot of developments are possible, for improving existing operators or for defining new classes of change operators.

Another possible way of development is to study the use of these belief change operators in other frameworks than classical logic. As illustrated by the works on horn clauses or on constraint networks, there are some subtleties that appear when one wants to work in these different frameworks.

We focused on purely qualitative approaches here, but there are also a lot of works on belief change (revision, update, merging, etc.) on quantitative frameworks. There are for instance a lot of works on ordinal conditional function [66, 67, 60], or on change of possibilistic logic bases [6, 7, 46, 8].

Merging is also at work on numerical datas, see for instance [65, 3, 10] for some examples of numerical data fusion. See [11] for an interesting global overview on (logical and numerical) merging.

References

1. C. Alchourrón and D. Makinson. On the logic of theory change: Safe contraction. *Studia Logica*, 44:405–422, 1985.
2. C. E. Alchourrón, P. Gärdenfors, and D. Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50:510–530, 1985.
3. A. Ayoun and P. Smets. Data association in multi-target detection using the transferable belief model. *International Journal of Intelligent Systems*, 16(10):1167–1182, 2001.
4. C. Baral, S. Kraus, and J. Minker. Combining multiple knowledge bases. *IEEE Transactions on Knowledge and Data Engineering*, 3(2):208–220, 1991.
5. C. Baral, S. Kraus, J. Minker, and V. S. Subrahmanian. Combining knowledge bases consisting of first-order theories. *Computational Intelligence*, 8(1):45–71, 1992.
6. S. Benferhat, D. Dubois, J. Lang, H. Prade, A. Saffiotti, and P. Smets. A general approach for inconsistency handling and merging information in prioritized knowledge bases. In *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, pages 466–477, 1998.
7. S. Benferhat, D. Dubois, and H. Prade. A computational model for belief change and fusing ordered belief bases. In H. Rott and M. A. Williams, editors, *Frontiers in Belief revision*. Kluwer, 1999.
8. S. Benferhat and S. Kaci. Fusion of possibilistic knowledge bases from a postulate point of view. *International Journal of Approximate Reasoning*, 33(3):255–285, 2003.
9. S. Benferhat, S. Lagrue, and O. Papini. Revision of partially ordered information : Axiomatization, semantics and iteration. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI'05)*, pages 376–381, 2005.
10. I. Bloch, T. Géraud, and H. Maître. Representation and fusion of heterogeneous fuzzy information in the 3d space for model-based structural recognition—application to 3d brain imaging. *Artificial Intelligence*, 148(1-2):141–175, 2003.
11. I. Bloch, A. Hunter, A. Appriou, A. Ayoun, S. Benferhat, P. Besnard, L. Cholvy, R. M. Cooke, F. Cuppens, D. Dubois, H. Fargier, M. Grabisch, R. Kruse, J. Lang, S. Moral, H. Prade, A. Saffiotti, P. Smets, and C. Sossai. Fusion: General concepts and characteristics. *International Journal of Intelligent Systems*, 16(10):1107–1134, 2001.

12. I. Bloch and J. Lang. Towards mathematical morpho-logics. In *Proceedings of the Eighth International Conference on Information Processing and Management of Uncertainty in Knowledge-Based Systems (IPMU'00)*, pages 1405–1412, 2000.
13. G. Boella, C. da Pereira, A. Tettamanzi, and L. van der Torre. Making others believe what they want. In *Artificial Intelligence in Theory and Practice II*, pages 215–224, 2008.
14. R. Booth. A negotiation-style framework for non-prioritised revision. In *Proceedings of the Eighth conference on Theoretical Aspects of Rationality and Knowledge (TARK'01)*, pages 137–150, 2001.
15. R. Booth. Social contraction and belief negotiation. In *Proceedings of the Eighth Conference on Principles of Knowledge Representation and Reasoning (KR'02)*, pages 374–384, 2002.
16. R. Booth. Social contraction and belief negotiation. *Information Fusion*, 7(1):19–34, 2006.
17. R. Booth and T. Meyer. Admissible and restrained revision. *Journal of Artificial Intelligence Research*, 26:127–151, 2006.
18. R. Booth, T. Meyer, and I. J. Varzinczak. Next steps in propositional horn contraction. In *Proceedings of the Twenty first International Joint Conference on Artificial Intelligence (IJCAI'09)*, pages 702–707, 2009.
19. R. Booth, T. Meyer, I. J. Varzinczak, and R. Wassermann. Horn belief change: A contraction core. In *Proceedings of the Nineteenth European Conference on Artificial Intelligence (ECAI'10)*, pages 1065–1066, 2010.
20. R. Booth and A. Nittka. Reconstructing an agent's epistemic state from observations. In *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence (IJCAI'05)*, pages 394–399, 2005.
21. J. F. Condotta, S. Kaci, P. Marquis, and N. Schwind. Merging qualitative constraint networks in a piecewise fashion. In *Proceedings of the Twenty first International Conference on Tools with Artificial Intelligence (ICTAI'09)*, pages 605–608, 2009.
22. J. F. Condotta, S. Kaci, P. Marquis, and N. Schwind. Merging qualitative constraints networks using propositional logic. In *Proceedings of the Tenth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'09)*, pages 347–358, 2009.
23. J. F. Condotta, S. Kaci, P. Marquis, and N. Schwind. Majority merging: from boolean spaces to affine spaces. In *Proceedings of the Nineteenth European Conference on Artificial Intelligence (ECAI'10)*, pages 627–632, 2010.
24. S. Coste-Marquis, C. Devred, S. Konieczny, M. C. Lagasque-Schiex, and P. Marquis. On the merging of dung's argumentation systems. *Artificial Intelligence*, 171:740–753, 2007.
25. M. Dalal. Investigations into a theory of knowledge base revision: preliminary report. In *Proceedings of the American National Conference on Artificial Intelligence (AAAI'88)*, pages 475–479, 1988.
26. A. Darwiche and J. Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, (89):1–29, 1997.
27. J. P. Delgrande. Horn clause belief change: Contraction functions. In *Proceedings of the Eleventh International Conference on the Principles of Knowledge Representation and Reasoning (KR' 08)*, pages 156–165, 2008.
28. J. P. Delgrande, D. Dubois, and J. Lang. Iterated revision as prioritized merging. In *Proceedings of the Tenth International Conference on Knowledge Representation and Reasoning (KR'06)*, pages 210–220, 2006.
29. J. P. Delgrande and R. Wassermann. Horn clause contraction functions: Belief set and belief base approaches. In *Proceedings of the Twelfth International Conference on the Principles of Knowledge Representation and Reasoning (KR' 10)*, 2010.
30. P. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77:321–357, 1995.

31. F. Dupin de Saint-Cyr and J. Lang. Belief extrapolation (or how to reason about observations and unpredicted change). *Proceedings of the Eighth Conference on Principles of Knowledge Representation and Reasoning (KR'02)*, pages 497–508, 2002.
32. F. Dupin de Saint-Cyr and J. Lang. Belief extrapolation (or how to reason about observations and unpredicted change). *Artificial Intelligence*, 175(2):760–790, 2011.
33. M. A. Falappa, G. Kern-Isberner, and G. R. Simari. Belief revision and argumentation theory. In Guillermo Simari and Iyad Rahwan, editors, *Argumentation in Artificial Intelligence*, pages 341–360. Springer US, 2009.
34. E. Fermé and S. O. Hansson, editors. *Journal of Philosophical Logic. Special Issue on 25 Years of AGM Theory*, volume 40, number 2. Springer Netherlands, 2011.
35. P. Gärdenfors. *Knowledge in flux*. MIT Press, 1988.
36. P. Gärdenfors, editor. *Belief Revision*. Cambridge University Press, 1992.
37. P. Gärdenfors and D. Makinson. Revisions of knowledge systems using epistemic entrenchment. In *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 83–95, 1988.
38. O. Gauwin, S. Konieczny, and P. Marquis. Conciliation and consensus in iterated belief merging. In *Proceedings of the Eighth European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'05)*, pages 514–526, 2005.
39. N. Gorogiannis and A. Hunter. Merging first-order knowledge using dilation operators. In *Proceedings of the Fifth International Symposium on Foundations of Information and Knowledge Systems (FoIKS'08)*, pages 132–150, 2008.
40. A. Grove. Two modellings for theory change. *Journal of Philosophical Logic*, 17(157-180), 1988.
41. S. O. Hansson. *A Textbook of Belief Dynamics*. Kluwer Academic Publishers, 1997.
42. S. O. Hansson, editor. *Theoria. Special Issue on non-prioritized belief revision*, volume 63, number 1-2. Wiley, 1997.
43. A. Herzig and O. Rifi. Propositional belief base update and minimal change. *Artificial Intelligence*, 115(1):107–138, 1999.
44. A. Hunter and J. P. Delgrande. Belief change in the context of fallible actions and observations. In *Proceedings of the Twenty first National Conference on Artificial Intelligence and the Eighteenth Innovative Applications of Artificial Intelligence Conference (AAAI'06)*, 2006.
45. Y. Jin and M. Thielscher. Iterated belief revision, revised. *Artificial Intelligence*, 171:1–18, 2007.
46. S. Kaci, S. Benferhat, D. Dubois, and H. Prade. A principled analysis of merging operations in possibilistic logic. In *Proceedings of the Sixteenth Conference in Uncertainty in Artificial Intelligence (UAI'00)*, pages 24–31, 2000.
47. H. Katsuno and A. O. Mendelzon. On the difference between updating a knowledge base and revising it. In *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning (KR'91)*, pages 387–394, 1991.
48. H. Katsuno and A. O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52:263–294, 1991.
49. H. Katsuno and A. O. Mendelzon. On the difference between updating a knowledge base and revising it. In P. Gärdenfors, editor, *Belief Revision*. Cambridge University Press, 1992.
50. S. Konieczny. Belief base merging as a game. *Journal of Applied Non-Classical Logics*, 14(3):275–294, 2004.
51. S. Konieczny, M. Medina Grespan, and R. Pino Pérez. Taxonomy of improvement operators and the problem of minimal change. In *Proceedings of the Twelfth International Conference on Principles of Knowledge Representation and Reasoning (KR'10)*, 2010.
52. S. Konieczny, J. Lang, and P. Marquis. DA^2 merging operators. *Artificial Intelligence*, 157(1-2):49–79, 2004.

53. S. Konieczny and R. Pino Pérez. Merging information under constraints: a logical framework. *Journal of Logic and Computation*, 12(5):773–808, 2002.
54. S. Konieczny and R. Pino Pérez. Confluence operators. In *Proceedings of the Eleventh European Conference on Logics in Artificial Intelligence (JELIA'08)*, pages 272–284, 2008.
55. S. Konieczny and R. Pino Pérez. Improvement operators. In *Proceedings of the Eleventh International Conference on Principles of Knowledge Representation and Reasoning (KR'08)*, pages 177–186, 2008.
56. P. Liberatore and M. Schaerf. Arbitration (or how to merge knowledge bases). *IEEE Transactions on Knowledge and Data Engineering*, 10(1):76–90, 1998.
57. P. Liberatore and M. Schaerf. BReLS: A system for the integration of knowledge bases. In *Proceedings of the Seventh Conference on Principles of Knowledge Representation and Reasoning (KR'00)*, pages 145–152, 2000.
58. J. Lin and A. O. Mendelzon. Merging databases under constraints. *International Journal of Cooperative Information System*, 7(1):55–76, 1998.
59. T. Meyer, N. Foo, D. Zhang, and R. Kwok. Logical foundations of negotiation: Strategies and preferences. In *Proceedings of the Ninth Conference on Principles of Knowledge Representation and Reasoning (KR'04)*, pages 311–318, 2004.
60. T. A. Meyer. On the semantics of combination operations. *Journal of Applied Non-Classical Logics*, 11(1-2):59–84, 2001.
61. P. Z. Revesz. Model-theoretic minimal change operators for constraint databases. In *Proceedings of the Sixth International Conference on Database Theory (ICDT'97)*, pages 447–460, 1997.
62. P. Z. Revesz. On the semantics of arbitration. *International Journal of Algebra and Computation*, 7(2):133–160, 1997.
63. N. D. Rotstein, M. O. Moguillansky, M. A. Falappa, A. J. García, and G. R. Simari. Argument theory change: Revision upon warrant. In *Proceeding of the 2008 conference on Computational Models of Argument (COMMA'08)*, pages 336–347, 2008.
64. H. Rott. *Change, choice and inference: a study of belief revision and nonmonotonic reasoning*. Oxford logic guides. Clarendon, 2001.
65. P. Smets. The combination of evidence in the transferable belief model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(5):447–458, 1990.
66. W. Spohn. Ordinal conditional functions: a dynamic theory of epistemic states. In W. L. Harper and B. Skyrms, editors, *Causation in Decision, Belief Change, and Statistics*, volume 2, pages 105–134. 1987.
67. M. A. Williams. Transmutations of knowledge systems. In *Proceedings of the Fourth International Conference on the Principles of Knowledge Representation and Reasoning (KR'94)*, pages 619–629, 1994.
68. D. Zhang, N. Foo, T. Meyer, and R. Kwok. Negotiation as mutual belief revision. In *Proceedings of the American National Conference on Artificial Intelligence (AAAI'04)*, pages 317–322, 2004.