A Visit to Belief Change without Compactness

Jandson S. Ribeiro
Department of Computer Science, University of São Paulo
São Paulo, Brazil
jandson@ime.usp.br
Department of Computing, Macquarie University
Sydney, Australia

Abstract

Dealing with dynamics is a vital problem in Artificial Intelligence (AI). An intelligent system should be able to perceive and interact with its environment to perform its tasks satisfactorily. To do so, it must sense external actions that might interfere with its tasks, demanding the agent to self-adapt to the environment dynamics. In AI, the field that studies how a rational agent should change its knowledge in order to respond to a new piece of information is known as Belief Change. It assumes that an agent's knowledge is specified in an underlying logic that satisfies some properties including compactness: if an information is entailed by a set X of formulae, then this information should also be entailed by a finite subset of X. Several logics with applications in AI, however, do not respect this property. This is the case of many temporal logics such as LTL and CTL. Extending Belief Change to these logics would provide ways to devise self-adaptive intelligent systems that could respond to change in real time. This is a big challenge in AI areas such as planning, and reasoning with sensing actions. Extending Belief Change beyond the classical spectrum has been shown to be a tough challenge, and existing approaches usually put some constraints upon the system, which are either too restrictive or dispense some of the so desired rational behaviour an intelligent system should present. This is a summary of the thesis "Belief Change without Compactness" by Jandson S Ribeiro. The thesis extends Belief Change to accommodate non-compact logics, keeping the rationality criteria and without imposing extra constraints. We provide complete new semantic perspectives for Belief Change by extending to non-compact logics its three main pillars: the AGM paradigm, the KM paradigm and Non-monotonic Reasoning.

Keywords: Belief Revision, Belief Update, Compactness, Non-Monotonic reasoning.

1 Introduction

Reasoning about dynamics is a challenge in both Computer Science (CS) and Artificial Intelligence (AI). Consider, for instance, the planning problem in AI: given a task, an agent needs to find a chain of actions (a plan) whose execution successfully achieves the task goal. Now imagine that an agent, say 'A', is executing a plan, and an external agent interferes with it (it blocks one of agent 'A' resources, for instance). In this case, agent 'A' original plan is no longer suitable, and the agent needs to come up with a whole new plan, considering the new circumstances of the environment. Completion of complex tasks, such as rescue missions, are susceptible to frequent external interferences. For instance, a malicious agent could attempt to steal information, or put a member of a rescue team on the wrong track. The nature of interference can be more innocuous: a piece of missing/unavailable information can interfere with the completion of the task. In these scenarios, the re-planning task should be performed automatically by the agent and it should guarantee the quality of its solution.

Software development is also dynamic, as a system often needs to be changed throughout its development process for many reasons: deal with a recently identified error, evolve the system to incorporate a new feature, or a change of requirements due to a misspecification. In this case, the system needs to be modified/repaired. Although, many approaches to automatically verify systems exist (e.g. Model Checking [1]), standard approaches to repair a system are mostly manual and exhaustive error-prone processes. Intelligent systems

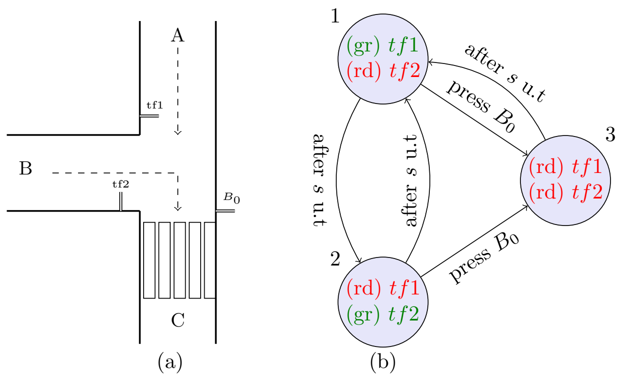

Figure 1: A crossing road (a) and an automatic traffic light system model (b). We use (gr) and (rd) to indicate respectively colours green and red, for instance at state 2, $(rd)\ tf1$ means $tf1$ is red. The u.t abbreviation means unit of time.

in AI could assist in handling dynamics in software development. For example, they could automate the process to modify/repair a system to accommodate new requirements or remove errors.

The core problem in reasoning about dynamics revolves around how an agent maintain its corpus of knowledge: any action an agent makes, including a response to an event (e.g. an external interference), is motivated and justified by the information the agent holds when the event is triggered. The field in AI that deals with the dynamics of an agent body of knowledge is known as Belief Change [2]. The two most influential theories in Belief Change are the AGM paradigm of *belief revision*, initially proposed in [3] and further developed in works such as [4], [5]; and the KM paradigm [6] of *belief update*. Both approaches conceptualise the notion of rational change motivated by the idea that the changes involved should be minimised. This minimality principle is formalised by rationality postulates which prescribe good behaviours for belief change operators. These paradigms also propose classes of belief change constructions strongly characterized by these rationality postulates.

This thesis considers the problem of how to deal with dynamics in AI, with the motivation that this can be extended to devise enhanced and autonomous intelligent systems. This thesis focuses on the foundations of Belief Change and AI. In Sections 1.2 and 5.1, we discuss in detail the impact and applications of this work.

The knowledge (or beliefs) of an agent is classically expressed via sentences specified in propositional logic. Relying solely on propositional logic to represent knowledge, however, is too restrictive. There are many other kinds of knowledge that need to be expressed in AI and computational systems. For instance, conceptual knowledge which can be expressed via Description Logics [7], Horn Logic [8] which are used for rule-based systems, and a variety of Modal Logics, such as Temporal Logics [9], which are widely used in planning and Formal Methods to specify the behaviour of the systems via temporal operators. To this end, considering Belief Change in the context of non-classical logics is essential to develop enhanced intelligent systems. Let us consider, for example, the problem of how a computer system could self-repair or repair other systems, by showing an intelligent system being developed in the context of smart cities. Figure 1 (a) illustrates a very simple scenario that this system should manage. There is a *"T" Junction* with two traffic lights, *tf1* and *tf2*, that respectively control the traffic flow to point C from points A and B respectively. A pedestrian can use the pedestrian crossing, Figure 1 (a), only when both the traffic lights, *tf1* and *tf2*, are red; and may press a button B_0 for this purpose.

A rather naïve implementation of this traffic control system is depicted in Figure 1 (b). This model illustrates the possible states of the system and how they transition from one state to another. By default, the system cycles between State 1 to State 2 after an elapse of s units of time. When button B_0 is pressed, the system goes to State 3, allowing pedestrians to cross the road, stays in that state for s units of time, and then transitions to State 1. At the first glance, this implementation of the traffic control system might appear to fit the bill.

Now, consider what is known as *liveness* – a very desirable property that a system is often required to satisfy. In the context of our traffic control system, a good instance of this property is the requirement that vehicles on both sides A and B, as well as pedestrians, should have the opportunity to go through the pedestrian crossing. As we will show shortly, the model of Figure 1 (b) does not satisfy liveness for the cars from direction B.

To see this, suppose that a pedestrian pushes the button B_0 to cross the road. This makes the system go to state 3. After s units of time elapses, the system goes to state 1 opening the traffic for cars from side A.

Then, before it changes to state 2, a new pedestrian requests to cross, moving the system to state 3 again. This loop of actions can continue indefinitely, preventing cars on side B to go through the crossing, thus violating the liveness property. Although such an indefinite line-up of pedestrians stretches the imagination, it demonstrates that the system is genuinely vulnerable to malicious attacks: instead of a huge number of pedestrians, it could be a single invader on the network that intends to disrupt the system. At the rush time, pedestrians continuously pressing button B_0 would significantly delay the cars on side B. Hence, there is a real need to fix the model in Figure 1 (b).

Standard approaches to repair a system consist in giving a counter-example (a piece of the model responsible for the failure) and letting the system designer manually repair the system. This makes the repair process an exhaustive task prone to introductions of errors. This simple model above is just a small piece of a complex system with possibly hundreds of states, and numerous constraints imposed upon them. Manually repairing such a system is impractical.

An ideal approach would not only identify causes of failure, but also suggest which modifications should be carried out in order to repair the model in order to satisfy the relevant requirements. In other words, we should *construct a system that is capable of repairing other systems.* However it is not obvious how such systems can be constructed. Existing proposals that suggest system repair at purely syntactical or structural levels are rather simplistic.

The system repair problem can actually be seen as an instance of the belief change problem: a system model is analogous to the knowledge of an agent, and the required specification that such a system has to satisfy corresponds to an acquired piece of information. The result of accommodating the new piece of information corresponds to the process of repairing the system. In this case, the agent changes its knowledge (respectively, changes the system model) in order to incorporate (satisfy) the new piece of information (requirement). From this perspective, belief change could lead to a rich notion of rational system repair.

Many approaches for formal specification and verification of systems rely on temporal logics. These logics are used in Formal Methods and many fields in AI, such as: planning, reasoning about actions, and sensing actions [10]. In these applications, an agent has to deal with dynamics, while its body of knowledge is usually represented via temporal logics. Although one can represent such kind of knowledge, dealing with the dynamics of temporal beliefs is still a tough task, and approaches to deal with changes are usually ad-hoc [10], [11].

Given the relevance of temporal logics in AI, understanding the dynamics involved in the management of knowledge in these logics will allow the construction of enhanced and robust intelligent systems. To this end, as for classical logics, a general framework capable of explaining and guiding belief change in these logics would be ideal. There are, however, a number of complicating factors in constructing this formal framework. The standard approaches of belief change make some strong assumptions about the background logic used to represent an agent's knowledge in terms of what properties they satisfy. Such assumptions may excessively limit when and how belief change can be performed. One of these assumptions is that the logic in question satisfies *compactness*: if a formula φ is entailed by a set of formulae K, then φ must also be entailed by a finite subset of K.

Example of non-compact logics include Hybrid logics (e.g. PDL[12]) which are used to represent the dynamics of computational systems; and more expressive temporal logics, such as CTL and LTL[1], which are among the most widely used logics in AI. As investigated by Ribeiro *et al.* [13], compactness has an important role in Belief Change theory, and dispensing with it has serious consequences. For instance, in the AGM paradigm the connection between the standard constructions of belief change and the rationality postulates is lost. Even worse, Guerra and Wassermann [14] have shown that it is not possible to define such constructions in logics without compactness. Regardless of how we intend to extend Belief Change to temporal logics, it is clear we have to deal with the problem of discarding compactness. This is the main purpose of this thesis: how to extend Belief Change for non-compact logics. Addressing compactness will not only allow us to apply belief change to a huge class of logics, but also to understand why the traditional mechanisms of belief change depend so strongly on this property and, therefore, to better understand belief change itself. In addition, it will allow us to assume very little about the underlying logic, and look at belief change more freely, focusing on the aspect of rationality change rather than on properties from the underlying logic. This text is a summary of some of the results present in the PhD thesis "Belief Change without Compactness", by Jandson S Ribeiro, submitted to University of São Paulo (USP) and Macquarie University under a cotutelle agreement. This summary is a visit to that thesis, and it does not present in detail all the discussed results.

1.1 Contributions

Belief Change is founded on three main pillars: (i) the AGM paradigm of belief revision, (ii) the KM paradigm of belief update and (iii) belief change as a non-monotonic system. The two main pillars, AGM and KM paradigms, have an extra-logical point of view on belief change: an agent's knowledge is represented via an underlying logic, whereas extra-logical constructions, such as epistemic preference relations, are used to devise belief change mechanisms. For the third pillar, an agent's knowledge is represented in an underlying logic whose consequence relation is non-monotonic: acquisition of new information may expel previous information contradictory with the new one. In this case, a natural question to pose is: which non-monotonic systems behave accordingly to the belief change rationality postulates? *This thesis extends all the three main pillars of Belief Change to the realm of non-compact logics. By doing so, we provide a complete redefinition of Belief Change to new fields.*

The first question to pose is if it is even possible to define belief change functions according to AGM and KM paradigms rationality postulates. Toward this end, we identify the properties a logic must possess in order to guarantee the existence of such functions, without depending on compactness. We do this for both the AGM paradigm and the KM paradigm. For the latter, we also identify the rationality postulates that emerge from the standard constructions of the KM paradigm in the absence of compactness. For the non-monotonic logics point of view, we define a new non-monotonic system that behaves as to the AGM rationality postulates of belief revision.

1.2 Significance and Impact

For decades, belief change approaches have relied on strong assumptions about the underlying logic used to express an agent corpus of beliefs. These assumptions, which include compactness, however, forbid applying belief change in expressive and interesting logical formalisms with applications in AI and CS. Negative results have shown that the classical semantic constructions and their connections with the rationality criteria cease in the absence of such assumptions. It is not clear, however, what is the role of these assumptions for the connection between the rationality postulates and their standard semantic constructions.

This thesis makes a big step towards answering this question: the rationality criteria are indeed consistent with non-compact logics, but the issue is that the classical constructions rely on these assumptions to semantically connect with the rationality postulates. *As one of reviewers of this thesis put: by extending Belief Change to accommodate non-compact logics, this thesis causes a so long needed paradigm disruption in Belief Change:* it provides the semantically counterparts of Belief Change without relying on strong assumptions, therefore making Belief Change more independent from the underlying logic. This disruption opens new research lines inside Belief Change itself, and provides new mechanisms to devise enhanced applications in both AI and CS. Consequently, this disruption compels a reinvestigation of the other forms of belief change (e.g. Iterative and Multiple Belief Change) in more expressive logics. We discuss in Section 5.1 some of these future lines of research.

1.3 Overview

In Section 1.4, we introduce some notation and technical background necessary to the understanding of this thesis. In Section 2, we briefly review the AGM paradigm. We then introduce the problem of AGM belief change in non-compact logics, and we show how to perform AGM belief change in logics that do not have the compactness property. We provide semantic constructions to the two main forms of AGM belief change: contraction and revision. We also show that the interdefinability of revision and contraction, via the Levi and Harper identities, does not depend on compactness. In Section 3, we review the KM paradigm and the consequences of dispensing with compactness in this system. We devise classes of (fully) KM rational functions, and we identify some issues with the original KM rationality postulates regarding non-finitary languages. In Section 4, we bring to light the connection between AGM revision and non-monotonic reasoning, and we devise a new non-monotonic system, without compactness, that reconnects with the AGM revision. In the final section, we summarise the main results obtained in this work and their significance. We end by discussing some research directions worth pursuing in future, and some relevant applications. Some of the results reported at Section 2 were previously published at [15]. Results presented at Section 3 and Section 4 were published at [16] and [17], respectively.

1.4 Notation and Technical Background

Given a set A, the power set of A will be denoted as 2^A. We use the terms formula and sentence interchangeably. We will use upper case Roman letters (A, B, X, Y ...) to denote sets, and lower case Greek letter $(\varphi, \psi, \alpha, \beta, \dots)$ will be used to denote formulae. Propositional symbols will be denoted by lower case Roman letters (p, q, r, ...). The maximal elements of a set A w.r.t a Partial Order $<$ are given by $max_<(A) = \{a \in A \mid \forall b \in A, a \not< b\}$.

The material implication will be represented as usual, that is, $\varphi \to \psi = \neg\varphi \vee \psi$. We consider a logic as a pair $\langle L, Cn\rangle$, where L is a language and Cn: $2^L \to 2^L$ is a logical consequence operator that maps a set of formulae to the set of all its inferred formulae. For readability, for any formula φ, the set $Cn(\{\varphi\})$ will often simply be written as $Cn(\varphi)$. We will often pretend that the consequence operation Cn itself represents a logic when no confusion is imminent. We limit ourselves to logics that are Tarskian, that is, logics whose consequence operator satisfies the following three properties:

1. **(Monotonicity)**: if $A \subseteq B$ then $Cn(A) \subseteq Cn(B)$;
2. **(Idempotence)**: $Cn(A) = Cn(Cn(A))$;
3. **(Inclusion)**: $A \subseteq Cn(A)$;

Apart from being Tarskian, the consequence operation is often granted some other properties in the AGM belief change literature, and they are often dubbed *AGM Assumptions*:

- **(deduction)**: $\varphi \in Cn(A \cup \{\psi\})$ iff $\psi \to \varphi \in Cn(A)$;
- **(supraclassicality)**: if φ is a logical consequence of A in classical propositional logic, then $\varphi \in Cn(A)$;
- **(compactness)**: if $\varphi \in Cn(A)$ then there is a finite subset A' of A such that $\varphi \in Cn(A')$.

For notational convenience, we take $A + \varphi$ to mean $Cn(K \cup \{\varphi\})$, where A is a set of formulae and φ a formula. The operator $+$ is in fact the *belief expansion* operator used to add a new piece of information φ to a set of formulae A, without taking into consideration whether or not φ is consistent with A. See Section 2 for a discussion about this operator.

A logic $\langle L, Cn\rangle$ is closed under classical negation iff the language L is closed under the negation operator: for each formula $\varphi \in L$, $Cn(\varphi) \cap Cn(\neg\varphi) = Cn(\emptyset)$, and $Cn(\{\varphi, \neg\varphi\}) = L$. In other words, the negation is interpreted classically. Analogously, the logic is closed under the disjunction if the language is closed under such a connective and it is also classically interpreted.

A theory or belief set is a set of formulae logically closed under the consequence operator Cn, that is, a theory A is a set of formulae such that $A = Cn(A)$. We will reserve the letter K to denote theories. Let $\mathcal{I}$ be an interpretation domain of a logic $\langle L, Cn\rangle$, and M a model in $\mathcal{I}$. Given a set of models A, $Th(A) = \{\varphi \mid \forall M \in A, M \models \varphi\}$ is the theory of the formulae satisfied by all models in A. Moreover, given a theory $X \subseteq L$, the set of models that satisfy all formulae in X is $[\![X]\!] = \{M \in I \mid \forall\varphi \in X, M \models \varphi\}$. For simplicity, given a set of formulae X and a model M, we will write $M \models X$ to mean that M satisfies every formulae in X. Furthermore, we will call a theory T a *complete theory* if and only if, for every sentence φ, either $\varphi \in T$ or $\neg\varphi \in T$, and use T_L to denote all such complete theories. Notice that, in logics closed under classical negation, the existence of complete theories is trivial.

A theory is said to be consistent if it does not entail some formula φ and its negation $\neg\varphi$. The set of all Consistent Complete Theories (CCT) that entail a formula φ is given by the set $\omega(\varphi) = \{T \in T_L \mid \varphi \in T\}$. Conversely the set of all consistent complete theories that do not entail a formula φ is given by $\overline{\omega}(\varphi) = \{T \in T_L \mid \varphi \notin T\}$. A consistent complete theory from $\overline{\omega}(\varphi)$ will be called a complement of φ.

2 AGM Paradigm Beyond Compactness

Portions of texts used in this section were borrowed from [15]. The AGM paradigm, initially proposed in [3] and further developed in works such as [2] and [5], is the most influential theory in the discipline of Belief Change. In the AGM paradigm, the beliefs of an agent is represented as sentences expressed via an underlying logic $\langle L, Cn\rangle$. An epistemic state comprises all the beliefs an agent hold in a given instant of time, and it is expressed in this paradigm as a theory. We recall from the introductory section that a theory is a set of formulae K logically closed under Cn, that is, $K = Cn(K)$. Besides being Tarskian, the underlying logic Cn is assumed to satisfy some properties: *deduction, supraclassicality* and *comapctness* see Section 1.4). Following [18] and [19], these conditions are here dubbed *AGM assumptions*.

An agent presents three actions regarding its epistemic states in the light of a formula φ:

Expansion: $K + \varphi$, add φ to K;

Contraction: $K \dot{-} \varphi$, relinquish φ;

Revision: $K * \varphi$, incorporate φ keeping consistency.

Expansion simply adds the new information to the current epistemic state of an agent: $K + \varphi = Cn(K \cup \{\varphi\})$. Contraction, on the other hand, removes the piece of information φ in hand. The revision operation incorporates the new piece of information and, unlike expansion, it maintains consistency of the new epistemic state (given that φ is consistent). The three operations are ruled by sets of rationality postulates inspired on the principle of minimal change: when changing its beliefs an agent should preserve most of its beliefs, modifying only those that are essential to accommodate or remove the piece of information in hand.

2.1 Contraction

In the AGM paradigm, the contraction form of belief change is governed by a set of six basic rationality postulate and two supplementary ones. Let $\mathbf{K}$ be the set of all theories, a belief change function is any function $f : \mathbf{K} \times L \rightarrow \mathbf{K}$ that maps a pair of epistemic state and formula to a new epistemic state. Let $\dot{-}$ be a belief change function, the AGM rationality postulates for contraction are:

$(\mathbf{K}_1^-)$	$K \dot{-} \varphi = Cn(K \dot{-} \varphi)$	(closure)
$(\mathbf{K}_2^-)$	$K \dot{-} \varphi \subseteq K$	(inclusion)
$(\mathbf{K}_3^-)$	If $\varphi \notin K$, then $K \dot{-} \varphi = K$	(vacuity)
$(\mathbf{K}_4^-)$	If $\varphi \notin Cn(\emptyset)$, then $\varphi \notin K \dot{-} \varphi$	(success)
$(\mathbf{K}_5^-)$	$K \subseteq (K \dot{-} \varphi) + \varphi$	(recovery)
$(\mathbf{K}_6^-)$	If $Cn(\varphi) = Cn(\psi)$, then $K \dot{-} \varphi = K \dot{-} \psi$	
$(\mathbf{K}_7^-)$	$(K \dot{-} \varphi) \cap (K - \psi) \subseteq K \dot{-} (\varphi \wedge \psi)$	
$(\mathbf{K}_8^-)$	If $\varphi \notin K \dot{-} (\varphi \wedge \psi)$ then $K \dot{-} (\varphi \wedge \psi) \subseteq K \dot{-} \varphi$	

The postulates (K_1^-) to (K_6^-) are basic rationality postulates, while (K_7^-) and (K_8^-) are the supplementary postulates and dictates rules to contract conjunctive formulae. For a discussion about the rationality behind these postulates, see [2], [5].

Any belief change operation that satisfies closure and success will be called a contraction operation. Moreover, any contraction operation that satisfy the six basic AGM rationality postulates of contraction will be called AGM rational, and if besides the six basic postulates it satisfies that two supplementary ones, then we will say that it is fully AGM rational.

The rationality postulates dictates what is a rational change, but it does not inform how to construct contraction operations that follows such guidelines. Many AGM rational contraction operations have been proposed in the literature, and the most influential one is known as *Partial Meet Function* [3]. In order to contract a formula φ from a theory K, a *partial meet function* resorts to *remainders*: maximal subsets of K that do not entail φ. A *partial meet function* then applies a strategy of select and intersect: it chooses some of the remainders of K modulo φ (this choice is rationalised by the idea that the formulae in the picked remainders are the most reliable beliefs of the agent), and then it intersects all the selected remainders with K, corresponding to the contraction result. If the selection is made according to a pre-order (a transitive and reflexive relation) on remainders, then the respective *partial meet function* is called a *transitive partial meet function.*

Theorem 1 ([3])**.** *In the presence of the AGM assumptions, a contraction function is AGM rational iff it is a partial meet function. Moreover, a contraction function is fully AGM rational iff it is a transitive relational partial meet function.*

2.2 AGM Contraction: Dispensing with Compactness

The representation results between partial meet functions and the AGM rationality postulates, presented in the previous section, depends strongly on the AGM assumptions. Negative results have shown that lifting compactness has serious consequences: Ribeiro *et al.* [13] has shown that the connection between partial meet functions and the AGM rationality postulates breaks down; even worse Guerra and Wassermann [14] has shown that it is not even possible to define partial meet in these logics. These negative results, however, bring more questions than answers. Precisely, what is the cause for this disruption between postulates and constructions in the absence of compactness? We show that the answer for this question is that the standard constructions depend on the compactness to establish its rational behaviours, but AGM rational constructions can be defined in non-compact logics. To show this, we need first to consider what is known as AGM-compliance: a logic Cn is said to be *AGM-Compliant* if and only if, with Cn as the background logic, it is possible to define a contraction operation that satisfies all the six contraction AGM postulates.

The first result of this thesis comprises AGM-compliance to a huge class of non-compact logics:

Theorem 2. *Every Tarskian logic closed under classical negation and disjunction is AGM- Compliant.*

2.3 Basic Contraction Rationality without Compactness

In this section, we devise new classes of AGM rational contraction functions. In the AGM approach, the partial meet contraction depends on remainder sets whose existence is guaranteed by the compactness property of the background logic. But since we do not have the compactness property to fall back upon, the contraction function we define will depend on a selection of complete consistent theories which will be intersected with the belief set K. Accordingly we assume a Choice Function (CF) $\delta : L \to 2^{T_L}$ that maps each formula φ of L to a set of complete theories $\delta(\varphi)$, subject to constraints:

1. $\delta(\varphi) \neq \emptyset$;
2. if $\varphi \notin Cn(\emptyset)$, then $\delta(\varphi) \subseteq \overline{\omega}(\varphi)$;
3. for any formulas φ and ψ, if $\varphi \equiv \psi$ then $\delta(\varphi) = \delta(\psi)$.

The purpose of a CF is to pick the best complete theories that do not entail a non-tautological formula φ (condition 2). A CF is not syntax-sensitive (condition 3). Condition 1 dictates that for any formula, at least one complete theory has to be selected. We define a new kind of contraction function:

Definition 3. *Let K be a theory, φ a formula, and δ a choice function. An operation $-_\delta$ is an Exhaustive Contraction Function (ECF) iff*

- *$K -_\delta \varphi = K \cap \bigcap \delta(\varphi)$, if $Cn(\emptyset) \subset K \cap Cn(\varphi) = Cn(\varphi)$ and either $\neg\varphi \notin Cn(\emptyset)$ or $\bot \in K$;*
- *$K -_\delta \varphi = K$, otherwise.*

An ECF works in the following way. A formula φ must only be retracted from a belief set K when: (1) K is not simply the set of all the tautological formulas, (2) φ is not a tautology and (3) φ is in K. These three constrains are jointly expressed as $Cn(\emptyset) \subset K \cap Cn(\varphi) = Cn(\varphi)$. Besides, φ has to be consistent, that is $\neg\varphi \notin Cn(\emptyset)$, in order to be retracted from K; or K is inconsistent. Otherwise, K is left untouched.

For illustration purposes, let us contract a formula from a theory expressed in a non-compact logic. The Linear Temporal Logic [1] is a good candidate for it. For simplicity, we will consider only two temporal operators of that logic: G and X. The former means *Globally (always) in the future*, and X means *in the "neXt" time instant*. We will keep the disjunction and negation of that logic which are interpreted classically. For our example, it will suffice to know two properties regarding these two operators. First, $Cn(\{Gp, p, Xp, X^2p, \ldots, X^np, \ldots\} \subseteq Cn(Gp)$. In other words, Gp implies that p is true in the current time instant and in all next future instants. As disjunction and negation are interpreted classically, note that formulas such as $Xp \to Gp$ also belong to $Cn(Gp)$. The second point we need to know is that $Cn(\{p, Xp, \ldots, X^np, \ldots\}) = Cn(Gp)$. For more details of this logic and its whole semantics, see [1].

Example 1. *Consider the theory $K = Cn(Gp)$ and we wish to contract Xp from it. Let our choice function be δ_1 where:*

1. *If $Cn(\psi) = Cn(Xp)$, then*
   $\delta_1(\psi) = \{Cn(\{p, \neg Xp, Xp \to Gp, \neg X^2p, (X^2p) \to Gp, \ldots\})\}$;
2. *Else, if $Cn(\psi) = Cn(\emptyset)$ then $\delta_1(\psi) = T_L$;*
3. *Else, $\delta_1(\psi) = \{S \in T_L \mid \psi \notin S\}$.*

If ψ is a tautology, we just let $\delta_1(\psi) = T_L$. The first constrain above regards the complete theories chosen to contract the formula Xp. From the semantics of LTL, it is easy to check that the only theory in $\delta_1(XP)$ is a complete consistent theory. As Xp is the only formula we are interested in retracting, for all other non-tautological formulas ψ we let δ_1 choose all the complete theories that do not imply ψ (third constrain). So, $K -_{\delta_1} Xp = K \cap \delta_1(Xp) = Cn(\{p, Xp \to Gp, X^2p \to Gp, \ldots\})$.

It is easy to notice that $-_{\delta_1}$ satisfies (K_1^-) to (K_4^-). Postulate (K_6^-) follows from condition 3 of the definition of CF. For Recovery (K_5^-), it suffices to note that $Xp \to Gp$ is in both $\delta_1(Xp)$ and K, whereby, $(K -_{\delta_1} Xp) + Xp = K$.

Now we reach our first representation result, Theorem 4 below: the class of the ECF comprises exactly all and only the AGM rational functions.

Theorem 4. *A contraction function satisfies (K_1^-) to (K_6^-) iff it is an ECF.*

2.4 Full Rationality

To semantically characterize all the eight AGM rationality postulates, we slightly modify the ECF to work over binary relations on CCTs. The main idea is that a preference relation $\leqslant$ reveals some hidden preference an agent has among its beliefs, and the choice function $\delta_\leqslant$ always choose the best CCTs modulo $\leqslant$. We call a choice function that backs upon a binary relation on CCTs a relational choice function. Not every binary relation, however, is suitable to represent this epistemic preference an agent has over its beliefs. We precisely identify two conditions that these binary relations should satisfy in order to yield fully AGM rational contraction functions:

$$\begin{tikzcd} A & B \\ C \arrow[u] \arrow[ur, dashed] & D \arrow[u] \arrow[ul, dashed] \end{tikzcd}$$

Figure 2: Mirroring. A is preferred to C, hence to D.

($\overline{\omega}$-Maximal Cut) for every non-tautological formula $\varphi \in L$, $\overline{\omega}(\varphi)$ has a maximal element w.r.t. $<$;

(Mirroring) if $S_1 \not\leq S_2$ and $S_2 \not\leq S_1$; then for any $S' \in T_L$, if $S_1 \leq S'$ then $S_2 \leq S'$.

The first condition on $<$, *$\overline{\omega}$-maximal cut*, is similar to the *Limit Assumption* of Lewis [20] and the *Finite Stopperedness* of Gärdenfors and Makinson [21]. It guarantees that for every formula φ, an agent chooses at least one complement theory of φ. The purpose of the $\overline{\omega}$-Maximal cut is to ensure that every formula to be dropped will be successfully relinquished, that is, the theory being contracted will be intersected by complements of φ. We call a binary relation that satisfies $\overline{\omega}$-Maximal-cut *contra-headed*. A relational choice function whose corresponding binary relation is contra-headed will be called an annulment. We will use μ instead of δ to denote annulment functions.

Definition 5. *Let $<$ be a contra-headed relation over T_L. An annulment is a function $\mu_< : L \to 2^{T_L}$ such that (i) $\mu_<(\varphi) = max_<(\overline{\omega}(\varphi))$ if $Cn(\varphi) \neq Cn(\emptyset)$; and (ii) $\mu_<(\varphi) =$ some $\emptyset \neq X \subseteq T_L$ otherwise.*

As for the second condition, *Mirroring* is similar to the modular relation defined in [22] which was based on *modular* partial orders of [23] and [24]. Though the concept of modular relation is confined to be a partial order, we impose no such restriction. The intuition behind mirroring is that if an agent has no preference between two theories A and B, then those that are preferable to A should also be preferable to B and vice versa. For instance, when dropping a formula φ an agent may choose among the four complements of φ: A, B, C and D. It prefers A to C and B to D, that is, its preference relation is $\{(C, A), (D, B)\}$ which is depicted in Figure 2 by solid arrows. So, it will choose both A and B to contract φ. However, there is no preference between C and D. According to mirroring, all theories that are preferable to C are also preferable to D (and vice versa). Thus, the pairs (C, B) and (D, A) also need to be present in the relation (depicted by dashed arrows in Figure 2).

We define now a special case of the ECF: the Blade Contraction Function (BCF, for short).

Definition 6. *A Blade Contraction Function (BCF) $\dot{-}_<$ is an ECF whose choice function is an annulment $\mu_<$ founded on a contra-headed $<$.*

The interesting aspect of BCFs is that we can study the rationality of the contraction functions that they yield by just looking at how these relations structure epistemic preferences. The $\overline{\omega}$-Maximal-cut guarantees postulate (K_7^-), while mirroring guarantees (K_8^-). This leads to our representation theorem for fully AGM rational functions on non-compact logics:

Theorem 7. *A contraction function is fully AGM rational iff it is a BCF whose contra-headed relation satisfies mirroring.*

2.5 AGM Revision without Compactness

The revision operation consists in adding a new information to a belief set guaranteeing its consistency if the new information is consistent. Let $*$ be a belief change operator, the AGM rationality postulates for revision are:

($\mathbf{K}_1^*$)	$K * \varphi = Cn(K * \varphi)$	(closure)
($\mathbf{K}_2^*$)	$\varphi \in K * \varphi$	(success)
($\mathbf{K}_3^*$)	$K * \varphi \subseteq K + \varphi$	(inclusion)
($\mathbf{K}_4^*$)	If $\neg\varphi \notin K$, then $K + \varphi \subseteq K * \varphi$	(preservation)
($\mathbf{K}_5^*$)	$K * \varphi = Cn(\bot)$ iff $\neg\varphi \in Cn(\emptyset)$	(consistency)
($\mathbf{K}_6^*$)	If $Cn(\varphi) = Cn(\psi)$, then $K * \varphi = K * \psi$	
($\mathbf{K}_7^*$)	$K * (\varphi \wedge \psi) \subseteq (K * \varphi) + \psi$	
($\mathbf{K}_8^*$)	If $\neg\psi \notin K * \varphi$, then $K * \varphi + \psi \subseteq K * (\varphi \wedge \psi)$	

The postulates (K_1^*) to (K_6^*) constitute the basic rationality postulates, while (K_7^*) and (K_8^*) are the supplementary postulates. For a discussion about the rationality behind these postulates, please refer to [2], [5].

The constructions proposed to give semantics to AGM revision, as for instance the Grove's System of Spheres [2], [25], strongly depends on the assumption of Compactness. In this section, similarly to how we did for contraction, we propose new classes (for both basic and full rationality) of AGM revision constructions without assuming compactness. Our approach to devise AGM revision functions relies, as for contraction, on a choice function. The difference here is that instead of choosing complements of a formula φ, a choice function shall choose among the Complete Consistent Theories that indeed entail φ. The reason for this is because a revision function wants to incorporate a formula φ rather than expelling it, thus instead of looking for complements, it will pick up supporters of φ. To avoid misinterpretation, a "choice" function $\delta : L \to 2^{T_L}$ for revision purposes will be called a casting function subject of the following constraints:

(1) $\delta(\varphi) \neq \emptyset$;

(2) if $Cn(\varphi) \neq Cn(\bot)$, then $\delta(\varphi) \subseteq \omega(\varphi)$;

(3) for every formulas φ and ψ, if $\varphi \equiv \psi$ then $\delta(\varphi) = \delta(\psi)$.

The only difference between the constraints put upon a casting function and those of a choice function (for contraction) is condition (2) which requires the casting function to choose complete theories that entail a formula φ. This restriction of picking only complete consistent theories that entail a formula φ is imposed only when φ is consistent. In the case that φ is inconsistent, this restriction is lifted, and δ can choose any complete consistent theories. Conditions (1) and (3) are respectively the same as in a choice function definition.

We define the Exhaustive Revision Functions (ERF, for short) which satisfy the six basic AGM revision rationality postulates. The intuition behind an ERF is that given a theory K and a formula φ: if either φ is inconsistent, or $K + \varphi$ is not inconsistent (this corresponds to K being consistent with φ), then such an expansion is the revision result. On the other hand, if φ is inconsistent with K, then the revision function can choose any theory that entails φ. Such a theory is constructed by the intersection of the theories provided by a casting function.

Definition 8. *Given a casting function δ, an Exhaustive Revision Function (ERF) $*_\delta$ is defined as*

1. $K *_\delta \varphi = K + \varphi$, *if* $\neg\varphi \notin K$ *or* $Cn(\varphi) = Cn(\bot)$;
2. $K *_\delta \varphi = \bigcap \delta(\varphi)$, *otherwise.*

The class of the ERFs comprises exactly the class of all the basic AGM rational revision function:

Theorem 9. *A revision function satisfies (K_1^*) – (K_6^*) iff it is an ERF.*

To semantically characterize all the eight AGM revision postulates, we constrain the class of ERFs in way that the choice function of each ERF picks its consistent complete theories according to a binary relation. All we need is *Mirroring* and:

(ω-**Maximal cut**) for every consistent formula $\varphi \in L$, $\omega(\varphi)$ has a maximal element w.r.t $<$.

We will refer to ω-maximal cut as simply maximal-cut. A casting function whose relation satisfies maximal-cut will be called a *coalition*, and a revision function constructed on a coalition function will be called a Constellar Revision Function. Maximal cut guarantees postulate (K_7^*), while mirroring captures (K_8^*). This give rise to our representation theorem for fully AGM rational revision functions:

Theorem 10. *A revision function is fully AGM rational iff it is a CRF $*_<$ such that $<$ satisfies mirroring.*

2.6 Levi and Harper Identities

Contraction and revision functions can be defined in terms of each other via the Levi and Harper identities:

(Levi Identity) $K * \varphi = (K \dot{-} \neg\varphi) + \varphi$
(Harper Identity) $K \dot{-} \varphi = K \cap (K * \neg\varphi)$.

This interdefinability is known to occur in the presence of the AGM assumptions which includes compactness. We show that Levi and Harper identities will only require a logic to be Tarskian and closed under classical negation and disjunction.

Theorem 11. *Given a contraction function $\dot{-}_\delta$ and the revision operator $*$ obtained from $\dot{-}_\delta$ via Levi Identity. If $\dot{-}_\delta$ satisfies* $(\boldsymbol{K_1^-})$ *to* $(\boldsymbol{K_4^-})$ *and* $(\boldsymbol{K_6^-})$*, then $*$ satisfies (K_1^*) to (K_6^*). Additionally, if $\dot{-}_\delta$ satisfies* $(\boldsymbol{K_1^-})$ *to* $(\boldsymbol{K_8^-})$ *then $*$ satisfies all eight AGM revision postulates.*

Theorem 12. *If a revision function $*$ satisfies* $(\boldsymbol{K_1^*})$*-*$(\boldsymbol{K_6^*})$ *(respectively (K_7^*) and (K_8^*)) then the contraction function $\dot{-}_\delta$ obtained via Harper identity satisfies* $(\boldsymbol{K_1^-})$ *to* $(\boldsymbol{K_6^-})$ *(respectively (K_7^-) and (K_8^-)).*

3 KM Paradigm

Portions of texts used in this section were borrowed from [17]. The AGM paradigm is undoubtedly the most influential theory in Belief Change literature. Another paradigm that is central in Belief Change is the KM paradigm, proposed by Katsuno and Mendelzon [6]. Both KM and AGM paradigms are seen as the two cornerstones that sustain Belief Change.

In the KM paradigm, dispensing with compactness also makes the connection between the standard constructions of belief update and the KM rationality postulates to break down. We show that in the KM case, the culprit is not compactness itself, but rather a property that non-compact logics have: non-finitaryness. A logic is non-finitary if it has an infinite number of models. It turns out that this poses as a serious problem that requires immediate solution, as besides for propositional logics and some of its fragments (such as Horn Logics), most of logics are non-finitary. The modal system K, for instance, which is a very simple and decidable normal modal logic, it is already non-finitary. We consider the problem of KM paradigm on non-finitary logics, specifically its update form of belief change: accommodate a piece of information and potentially keep consistency of an agent body of knowledge.

Originally, the KM paradigm represents an agent belief state as a single finite formula. This is because, every theory can be finitely represented in finitary logics. When dealing with non-finitary logics, however, representing epistemic states in this way may be inconvenient. Mainly because in these logics not every belief set can be represented via a single formula or a finite set of formulae. We translate the rationality postulates and the standard constructions of update considering epistemic states as theories.

(K⋄1) $K \diamond \varphi$ is a theory;
(K⋄2) $\varphi \in K \diamond \varphi$;
(K⋄3) If $\varphi \in K$, then $K \diamond \varphi = K$;
(K⋄4) $K \diamond \varphi$ is consistent, if both K and φ are consistent;
(K⋄5) If $Cn(\varphi) = Cn(\psi)$, then $K \diamond \varphi = K \diamond \psi$;
(K⋄6) $K \diamond (\varphi \wedge \psi) \subseteq K \diamond \varphi + \psi$;
(K⋄7) If $\psi \in K \diamond \varphi$ and $\varphi \in K \diamond \psi$, then
$K \diamond \varphi = K \diamond \psi$;
(K⋄8) If K is complete, then
$K \diamond \varphi \vee \psi \subseteq (K \diamond \varphi) \cup (K \diamond \psi)$;
(K⋄9) $(K \cap K') \diamond \varphi = (K \diamond \varphi) \cap (K' \diamond \varphi)$.

As we require an epistemic state to be a theory, the postulate (K⋄1) is added here with the purpose to ensure that every update function goes from a belief set to another one. For a discussion about the rationality behind the remaining postulates, please refer to [6]. The rationality postulates give rise to a class of belief change functions, here dubbed the Classical Update (CUP) functions

$$\text{CUP}: \quad K \diamond \varphi = Th\Big(\bigcup_{M \in [\![K]\!]} Max_{\leqslant_M}([\![\varphi]\!])\Big).$$

Given a theory K and a formula φ, a CUP operator $\diamond$ behaves as follows. For each model M of K, the operator $\diamond$ chooses from the models of φ those ones closest to M. The notion of distance between models is usually given by a pre-order $\leqslant_M$. A model M_1 is closer to M than a model M_2, if $M_2 \leqslant_M M_1$. These selected models are then assembled, and the theory of these models corresponds to $K \diamond \varphi$. Each model M is assigned to a pre-order $\leqslant_M$ subject to:

faithfulness if $M' \neq M$, then $M' <_M M$, for every model M'.

A function that assigns each model M to a partial pre-order $<_M$ that satisfies the above constrains is called a *faithful* assignment. The intuition behind the *faithfulness* constrain is that a relation $<_M$ assigned to a model M puts M as the most preferable model of that relation. Katsuno and Mendelzon [6] show that in Propositional Logic, the KM update postulates characterize the CUP functions:

Theorem 13. *A belief change function satisfies the KM update postulates iff it is a CUP function.*

The connection between the KM update postulates and CUP functions cease for non-finitary logics. For instance, for non-finitary logics CUP functions do not satisfy postulate (K$\diamond$7). The reason for this is related to the characteristic of non-finitary logics having an infinity number of models. In this case, a pre-order may induce an infinite chain in which it is not possible to determine a maximal (resp. minimal) element.

We will propose, in the next section, a new class of functions that satisfy all the KM update postulates.

3.1 KM Basic Rationality

The functions we will devise here will operate over complete consistent theories rather than models, following our approach in Section 2.3. In this way, we bring an uniformity between the semantics constructions of AGM and KM paradigms. We propose here to split the update rationality postulates into two groups: the basic rationality postulates which comprises postulates (K$\diamond$1) to (K$\diamond$5) and (K$\diamond$9); and the supplementary postulates (K$\diamond$6) to (K$\diamond$8). Any belief change function that satisfies all the basic update postulates is dubbed an update function.

We will assume that, instead of a pre-order, an agent has an appointee function that judges which complete theories of a formula φ are closest to a complete theory K. This appointee function is used to perform the update of K by φ. We recall, from the Introductory section that T_L stands for the set of all complete theories.

Definition 14. *An appointee is a function $\delta : T_L \times L \to 2^{T_L}$ that maps each consistent complete theory and formula φ to a set of consistent complete theories subject to:*

(D1) $\delta(S, \varphi) \neq \emptyset$*, if* S *and* φ *are consistent;*

(D2) $\delta(S, \varphi) \subseteq \omega(\varphi)$*;*

(D3) $\delta(S, \varphi) = \{S\}$*, if* $\varphi \in S$*;*

(D4) $\delta(S, \varphi) = \delta(S, \psi)$*, if* $Cn(\varphi) = Cn(\psi)$*.*

The purpose of an appointee is to select the best complete consistent theories that satisfy a formula φ (condition D2). This criterion is given locally, and it depends on the fixed complete theory S as background. An appointee is compelled to choose at least one consistent complete theory that contains φ, as long as φ is consistent (condition D1). If a formula φ belongs to a complete theory S, then the best theory for φ is S itself (condition D3). Finally, condition (D4) simply determines that an appointee is syntax independent.

We define a new class of update functions: the splinter functions which is similar in spirit to CUP functions.

Definition 15. *Given an appointee δ, a splinter is a function $\diamond_\delta$ such that*

1. *if K and φ are consistent, then*
$$K \diamond_\delta \varphi = \bigcap_{S \in \omega(K)} \big(\bigcap \delta(S, \varphi)\big);$$

2. *otherwise, $K \diamond_\delta \varphi = Cn(\bot)$.*

We recall from the introduction section that $\omega(K)$ corresponds to the set of all complete consistent theories that contains a theory K. The splinter functions semantically characterise the the basic KM update postulates: (K$\diamond$1) to (K$\diamond$5) and (K$\diamond$9).

Theorem 16. *A belief change function satisfies postulates (K$\diamond$1) to (K$\diamond$5) and (K$\diamond$9) iff it is a splinter.*

3.2 Revisiting the Supplementary Postulates

To capture the supplementary postulate, we make the appointees resort to binary relations in order to pick consistent complete theories. An appointee δ will assign, for each consistent complete theory K, a binary relation $\leqslant_K$. The appointee chooses the maximal elements within $\omega(\varphi)$ modulo $\leqslant_K$. We impose these relations to satisfy maximal-cut, which was introduced in Section 2.5, with the purpose to guarantee that for every formula φ, as long as it is consistent, the agent chooses at least one consistent complete theory of φ. This ensures that φ will be successfully incorporated through the update process. *Appointee functions that rely strictly on relations that satisfy maximal-cut will be called relational appointees.*

Relational appointees founded on maximal-cut alone are strong enough to yield splinters that satisfy postulates (K$\diamond$6) and (K$\diamond$8). However, postulate (K$\diamond$7) is still not enforced. For this, we will need Quasi-reflection.

Quasi-reflection: if $A \nleqslant B$ and $B \nleqslant A$ but $A \leqslant C$ then either

(i) $B \leqslant C$; or

(ii) if $C \leqslant C'$, then $B \leqslant C'$.

Quasi-reflection behaves similarly to mirroring. For each two elements A and B, there are two possibilities. If there is no preference between A and B, then A and B can mirror each other preferences, in this case it behaves exactly like mirroring. For the second condition, instead of making A and B mirror each other preferences, quasi-reflection allows them to skip the mirroring process one level above. So, for instance, if $B \leqslant C$ and $C \leqslant D$, then instead of enforcing $A \leqslant C$, as condition (i) demands, condition (ii) allows A to mimic one of the elements immediately above B, that is, mimicking C preferences. Thus, making $A \leqslant D$.

If we resort to the class of royal splinters whose appointees satisfy quasi-reflection (a quasi-reflected appointee), then we capture postulate (K$\diamond$7). We define then a new class of splinters, the royal splinters: *a splinter $\diamond_\mu$ is royal iff its appointee μ is relational and each assigned relation satisfies quasi-reflection.* We establish a representation theorem between the class of royal splinter founded on quasi-reflected appointees:

Theorem 17. *An update function is fully KM rational iff it is a royal splinter.*

3.3 Do We Need Total Relations?

The following postulate was proposed to be used in place of postulates (K$\diamond$7) and (K$\diamond$8):

(K$\diamond$10) if K is complete and $\neg\psi \notin K \diamond \varphi$ then
$K \diamond \varphi + \psi \subseteq K \diamond \varphi \wedge \psi$.

Postulate (K$\diamond$10) actually corresponds to postulate (K_8^*) of AGM revision. (see Section 2.5), with the restriction that it applies only to complete theories. Katsuno and Mendelzon [6] have shown that the set of postulates that results from replacing (K$\diamond$7) and (K$\diamond$8) for (K$\diamond$10) characterises exactly the class of CUP functions defined over total pre-orders. Use of total pre-orders to represent preference relations on beliefs has received many criticisms [4], [26]. We show that, actually, we do not need total pre-orders to capture (K$\diamond$10). All we need is to restrict to royal splinters whose appointees are founded on relations satisfying mirroring (introduced in Section 2.4). We shall refer to these royal-splinter as mirrored royal splinters. This leads us to the following representation result:

Theorem 18. *Every mirrored royal splinter satisfies (K$\diamond$10). Every update function that satisfies (K$\diamond$6) and (K$\diamond$10) is a mirrored royal splinter.*

4 Non-Monotonic Reasoning and Belief Change

Portions of texts used in this section were borrowed from [16]. Classical logics are monotonic: accumulation of new information (in form of additional premises) does not invalidate old conclusions. Commonsense reasoning, however is non-monotonic: acquisition of new information may "deactivate" some premisses that conflict with the new piece of information. When viewed from this angle, non-monotonic reasoning appears to be closely connected to accounts of rational belief change, such as the classic AGM account. Indeed it has been argued that belief dynamics and non-monotonic reasoning are different perspectives on the same phenomenon [27]. This idea is concisely captured in the standard notation:

BRNM: $x \mid\!\sim_K y$ if and only if $y \in K * x$

where K represents a fixed background knowledge, $\mid\!\sim_K$ is the non-monotonic inference operation that employs K in the background, and $*$ is a belief *revision* operation that yields the new "belief set" $K * x$ from the old belief set K in light of evidential input x. Oftentimes the subscript $_K$ from the relation $\mid\!\sim_K$ is dropped for notational convenience when the intention is clear from the context.

The connection between belief revision and non-monotonic reasoning captured by BRNM comes out handy in going back and forth between these two systems. Since non-monotonic reasoning and belief revision are strongly connected, the nature of non-monotonic reasoning is worth enquiring when the background logic of the corresponding belief revision is not assumed to be compact. We show that the connection from non-monotonic systems to belief revision cannot be completed in absence of the compactness assumption. We then devise a new non-monotonc system that reconnects with the AGM revision without assuming compactness.

4.1 Non-Monotonic Reasoning

In the non-monotonic reasoning system based on expectation orderings there are seven basic axioms and two supplementary axioms. This separation into two groups, first proposed in [21], makes the alignment of these axioms with the postulates for AGM belief revision explicit.

We will assume that a non-monotonic inference relation, denoted $\mid\!\sim$, builds upon an underlying logic Cn. We will only assume that Cn is Tarskian and closed under classical negation and disjunction. We require no other assumptions for the results of this section. The postulates (or axioms) for non-monotonic inference relation $\mid\!\sim$ are:

N1 If $\psi \in Cn(\varphi)$, then $\varphi \mid\!\sim \psi$
N2 If $Cn(\varphi) = Cn(\psi)$ and $\varphi \mid\!\sim \alpha$ then $\psi \mid\!\sim \alpha$
N3 If $\varphi \mid\!\sim \psi$ and $\alpha \in Cn(\psi)$ then $\varphi \mid\!\sim \alpha$
N4 If $\varphi \mid\!\sim \psi$ and $\varphi \mid\!\sim \alpha$ then $\varphi \mid\!\sim \psi \wedge \alpha$
N5 If $\varphi \mid\!\sim \psi$ then $\mid\!\sim \varphi \rightarrow \psi$
N6 If $\not\mid\!\sim \neg\varphi$ and $\mid\!\sim \varphi \rightarrow \psi$, then $\varphi \mid\!\sim \psi$
N7 If $\varphi \mid\!\sim \perp$ then $Cn(\varphi) = Cn(\perp)$
N8 If $\varphi \wedge \psi \mid\!\sim \phi$ then $\varphi \mid\!\sim \psi \rightarrow \phi$
N9 If $\varphi \not\mid\!\sim \neg\psi$ and $\varphi \mid\!\sim \phi$ then $\varphi \wedge \psi \mid\!\sim \phi$

Postulates **N1** to **N7** are dubbed the basic postulates, while **N8** and **N9** are dubbed the supplementary postulates. For simplicity, we will call any inference relation that satisfies the basic postulates a *non-monotonic inference relation*.

When the underlying logic Cn satisfies compactness, axioms N3 and N4 jointly imply:

Closure$_{\text{NM}}$ If $\varphi \mid\!\sim \psi$ for all $\psi \in A$, then $\varphi \mid\!\sim \alpha$ for all $\alpha \in Cn(A)$,

named after the closure postulate $(\mathbf{K}_1^*)$ of belief revision that it corresponds to [21]. On the other hand, postulates **N5** through **N7** respectively correspond to the AGM revision postulates $(\mathbf{K}_3^*)$ through $(\mathbf{K}_5^*)$. Furthermore, postulate **N2** corresponds to $(\mathbf{K}_6^*)$. Similarly, postulates **N8** and **N9** correspond respectively to the revision postulates $(\mathbf{K}_7^*)$ and $(\mathbf{K}_8^*)$ - see Section 2.5 for a complete list of the postulates.

This nice-correspondence between revision and non-monotonic inference relation $\mid\!\sim$ breaks down in the absence of compactness in many points. We start by stressing that without compactness **Closure**$_{\text{NM}}$ no longer follows from **N3** and **N4**:

Theorem 19. *If the underlying logic Cn is not compact, then **N3** and **N4** do not imply **Closure**$_{\text{NM}}$.*

Let us at this point make the following observation, that if we start with a revision operator $*$, obtain the inference $\mid\!\sim$ induced by it, and then again obtain the revision operator induced by this inference relation in turn, we will get back the revision operator $*$ we started with. Similarly, if we started with an inference operation $\mid\!\sim$ and obtain another inference via a revision operation induced by it, we will go back to our original inference operation $\mid\!\sim$.

The first observation that we make is that even that in the absence of compactness, as long as $*$ satisfies the six basic (respectively the supplementary) AGM revision postulates, the induced $\mid\!\sim$ relation satisfies the seven basic (respectively the supplementary) axioms of non-monotonicity.

Theorem 20. *If a belief change operator $*$ is (basic) AGM rational then its induced inference relation $\mid\!\sim$ is a non-monotonic inference relation. Additionally,*

*(i) if $*$ satisfies $(\mathbf{K}_7^*)$, then $\mid\!\sim$ satisfies **N8**;*

*(ii) if $*$ satisfies $(\mathbf{K}_8^*)$, then $\mid\!\sim$ satisfies **N9**.*

Although AGM rational function yelds non-monotonic inference relations satisfying the basic (respectively supplementary) axioms, the converse path breaks in many points. We exploit, in the remaining of this section, these issues; and in the next section we show how to restore the bridge between NMR and AGM revision.

Let us first show that the basic non-monotonic axioms do not correspond to the basic AGM revision postulates. For this purpose we need to construct a non-monotonic inference relation $\mid\!\sim$ such that the revision function $*$ induced by it violates some of the basic AGM postulates. We will conveniently take a revision function $*$ which violates one of the AGM postulates (namely, $(\mathbf{K}_3^*)$), and then we show that the non-monotonic inference relation $\mid\!\sim$ induced by it satisfies *all* the basic non-monotonic axioms. We construct such a belief revision operator in Example 2.

Example 2. *Let $\otimes$ be a fully AGM rational revision function, and p an arbitrary formula. The belief revision operation $*$ is constructed as: (i) $K * \varphi = (K + \varphi) + p$, if $\varphi \rightarrow \neg p \notin K$; and (ii) $K \otimes \varphi$, otherwise.*

The revision operation from Example 2 behaves in a simple way. It adds both p and a formula φ to K, if both φ and p are jointly consistent with K. On the other hand, if K is inconsistent with φ or p, it resorts to the rational AGM function $\otimes$. It is trivial to show that although it violates postulate $(\mathbf{K}_3^*)$, the inference relation $\mid\!\sim$ induced by it satisfies **N1** to **N9**. This example lead us to the following result:

Theorem 21. *There is a non-monotonic inference relation that satisfies all the non-monotonic axioms, but violate some of the basic AGM revision postulates.*

4.2 Bridging the Gap

In this section, we restore the bridge between NMR and AGM revision by introducing a new Non-Monotonic system.

There are precisely three AGM postulates which are not captured by the NMR axioms alone: $(\mathbf{K}_1^*)$, $(\mathbf{K}_3^*)$ and $(\mathbf{K}_4^*)$. The first one we address is $(\mathbf{K}_1^*)$ (the closure).

As we discussed in Section 4.1, when the underlying logic Cn is compact, axioms **N3** and **N4** trivially imply $\mathbf{Closure}_{\text{NM}}$. However, when Cn is not assumed to be compact, **N3** and **N4** may no longer guarantee $\mathbf{Closure}_{\text{NM}}$ (Theorem 19). Consequently, a non-monotonic inference may fail to induce a belief change operator $*$ that satisfies $(\mathbf{K}_1^*)$. Thus, satisfaction of $(\mathbf{K}_1^*)$ necessitates the presence of the $\mathbf{Closure}_{\text{NM}}$ axiom, while **N3** and **N4** can be replaced by $\mathbf{Closure}_{\text{NM}}$:

Proposition 22. *An inference relation satisfies* $\boldsymbol{Closure}_{\text{NM}}$ *if and only if its induced belief change operator satisfies* $(\boldsymbol{K}_1^*)$*. Moreover,* $\boldsymbol{Closure}_{\text{NM}}$ *implies* ***N3*** *and* ***N4****.*

We still have two postulates to address: (K_3^*) and (K_4^*). We showed in Example 2 that non-monotonic inference relations may not capture postulate $(\mathbf{K}_3^*)$. The example explored the fact that a tautological evidence allowed the inference of new information not previously believed: recall that $Cn(\top) * \top$ entailed the formula p which is clearly not in $Cn(\top)$.

This is inconsistent with postulate (K_3^*), which imposes that in reasoning in the light of a tautology, no new information shall be acquired. To workaround this problem, we introduce the following condition that guarantees postulate $(\mathbf{K}_3^*)$:

(Keeper) if K is consistent and $\mid\!\sim_K \varphi$ then $\varphi \in K$.

Proposition 23. *If an inference relation* $\mid\!\sim$ *satisfies both* ***N5*** *and the* ***Keeper****, then its induced belief change operator* $*$ *satisfies* $(\boldsymbol{K}_3^*)$*.*

Though **Keeper** guarantees satisfaction of (K_3^*), it does not capture postulate (K_4^*). This is because the role of **Keeper** is to forbid the insertion of spurious information, but it does not prevent the loss of relevant information, which is exactly the role of postulate (K_4^*).

Proposition 24. *There is one non-monotonic inference relation* $\mid\!\sim$ *that satisfies all the basic axioms of non-monotonicity and the Keeper, but its induced belief change operator does not satisfy* (K_4^*)*.*

To avoid the loss of relevant information, we introduce the following condition:

(Rooting) if K is consistent and $\varphi \in K$ then $\mid\!\sim_K \varphi$.

The **Rooting** axiom is the converse of the **Keeper** and enforces that the formulae present in a theory K should remain in the non-monotonic inferences by a tautology. The *Rooting* together with the *Keeper* and **N6** captures $(\mathbf{K}_4^*)$.

Proposition 25. *If an inference relation satisfies Keeper, Rooting and* ***N6****, then its induced belief change operator satisfies* $(\boldsymbol{K}_4^*)$*.*

Now we have all the pieces of the puzzle to show the first result of the representation theorem:

Theorem 26. *Let* $\mid\!\sim$ *be a non-monotonic inference relation, if it satisfies the Keeper, Rooting and* $Closure_{\text{NM}}$ *then its induced belief revision operator* $*$ *is AGM rational. Moreover, an AGM rational belief revision operator induces a non-monotonic inference that satisfies* ***Keeper****,* ***Rooting*** *and* $\boldsymbol{Closure}_{\text{NM}}$*.*

We reach our representation theorem, in the form of Corollary 27 below, which follows from Theorems 26 and Proposition 22. This result establishes a bridge between the basic AGM belief revision postulates with a new non-monotonic system that comprises the following axioms: N1, N2, N5-N7, **Keeper**, **Rooting** and $\mathbf{Closure}_{\text{NM}}$.

Corollary 27. *A belief change operator* $*$ *is AGM rational iff its induced inference relation* $\mid\!\sim$ *satisfies* ***N1****,* ***N2****,* ***(N5 - N7)****,* ***Keeper****,* ***Rooting****,* $\boldsymbol{Closure}_{\text{NM}}$*. If* $*$ *is fully AGM rational iff it satisfies* ***N1****,* ***N2****,* ***(N5 - N9)****,* ***Keeper****,* ***Rooting****,* $\boldsymbol{Closure}_{\text{NM}}$*.*

5 Concluding Remarks

AGM Paradigm sans Compactness

We have identified a huge class of AGM-compliant logics that do not depend on compactness: those closed under classical negation and disjunction. We then, provided novel constructions for AGM contraction and revision in these logics. Our constructions require only two conditions: *mirroring* and *maximal-cut*. Besides being more general than the standard constructions, as it assume very little about the underlying logic, these conditions allow the specification of preference relation that are not total. Totality has been highly criticized in the literature. Moreover, we also showed that the interdefinibility between contraction and revision in these logics via Levi's and Harper's identities do not depend on compactness.

KM Update in Non-Finitary Logics

We have considered the KM paradigm of belief update form of belief change, and showed that the connection between the standard constructions and the rationality postulates also breaks down in the absence of compactness. We have identified that the culprit behind this break-down is non-finitereness, a property present in logics without compactness. As a remedial measure, we devised a novel class of belief update constructions that reconnects with the KM rationality postulates without assuming neither compactness nor finitereness. We only require the underlying logic to be closed under classical negation and disjunction. Our approach is based on binary relations that satisfy two properties: *maximal-cut* and *quasi-reflection*. We then turn to the issue of total relations to represent the epistemic preferences of an agent over its beliefs, which has been highly criticized in the Belief Change literature. We show that, in the KM update, totality can be avoid by using *mirroring* instead.

AGM Paradigm and Non-monotonic Reasoning

Makinson and Gärdenfors [27] has shown that when the underlying logic is classical, non-monotonic systems based on expectation (NME, for short) and AGM revision are equivalent. This correspondence, however, depends on the standard AGM constructions which are not AGM rational in the absence of compactness. On the negative side, we showed that the connection between the NME and the AGM revision breaks down in the absence of compactness. On the positive side, we were able to restore this bridge by defining a new non-monotonic system. In particular, we have introduced two new axioms to NME. An interesting aspect of our result is that it does not depend on a specific class of AGM rational constructions. Instead, our proofs make a direct connection between the AGM revision postulates and the new non-monotonic system axioms. This is important, since it helps to understand how these two systems behave independently of extra-logical constructions.

5.1 Impact and Future Works

We discuss some directly impact of this thesis, as well as some immediate questions worth to address in a near future.

5.1.1 Automatic System Repair

Belief change mechanisms could be used to guide automatic system repair via its rationality postulates and semantic constructions. Repairing a system specification to satisfy an imposed constraint can be seen as the belief revision/update process of accommodating a new piece of information (imposed constraint) to the agent epistemic state (system specification). Many logics used to specify system behaviours formally are not compact (e.g. LTL, CTL and μ-calculus [1]). As we extended Belief Change to non-compact logics, automatic system repair guided by rational change becomes possible. The semantic constructions provided in this thesis can be used to give support in the development of repairing formal system specification.

5.1.2 Multiple Belief Change in Non-Finitary Languages

The case when an agent has to change its own beliefs in light of a set of sentences is known as *Multiple Belief Change* (MBC), see [28], [29]. This generalization is not so straightforward, and it might be even more difficult when the input set of sentences is infinite, known as Infinitary Belief Change [29].

In the case of non-compact and non-finitary logics, MBC and the special case of *Infinitary Belief Change* emerge as big challenges. Investigating how BCF can be extended to both non-compact logics and non-finitary logics emerges as an important future line of research. The constructions and results we presented here shall underpin the investigations for MBF in such logics.

5.1.3 Non-Compact Dynamic Logics

Dynamic logics such as Dynamic Epistemic Logics (DEL) and Dynamic Doxastic Logics (DDL) are used to reason about the dynamics involved in changing an agent's belief/knowledge. Whereas AGM and KM paradigms conceptualise the dynamism as an extra logical feature; DDL and DEL capture these dynamics in the object language via modal operators rationalized by the AGM/KM postulates. These Dynamic Logics conceptualise Knowledge/Beliefs as information expressed in compact logics. This represents a severe limitation when we look at some sub-disciplines in AI, such as planning, where conceptualising and modelling temporal knowledge is fundamental. The results of this thesis can be used to provide a new semantics for Dynamic logics in the context of non-compact logics.

Acknowledgement

This thesis were supervised by both Prof Renata Wassermann (University of São Paulo, Brazil) and Prof Abhaya Nayak (Macquarie University, Australia). I acknowledge the financial support from both CAPES (a Brazilian research foundation) for a scholarship which supported my work during my stay at University of São Paulo; and Macquarie University for a scholarship that helped me to conduct my PhD work during my stay in Sydney.

References

[1] E. M. Clarke, O. Grumberg, and D. A. Peled, *Model checking.* MIT Press, 2001.

[2] P. Gärdenfors, *Knowledge in flux: Modeling the dynamics of epistemic states.* The MIT press, 1988.

[3] C. E. Alchourrón, P. Gärdenfors, and D. Makinson, "On the logic of theory change: Partial meet contraction and revision functions," *The Journal of Symbolic Logic*, vol. 50, no. 2, pp. 510–530, 1985. DOI: 10.2307/2274239.

[4] H. Rott, "Preferential belief change using generalized epistemic entrenchment," *Journal of Logic, Language and Information*, vol. 1, no. 1, pp. 45–78, 1992. DOI: 10.1007/BF00203386.

[5] S. O. Hansson, *A Textbook of Belief Dynamics: Theory Change and Database Updating.* Kluwer Academic, 1999, vol. 11.

[6] H. Katsuno and A. O. Mendelzon, "On the difference between updating a knowledge base and revising it," in *Belief revision*, P. Gärdenfors, Ed., Cambridge University Press, 2003, pp. 183–203.

[7] F. Baader, D. Calvanese, D. McGuinness, P. Patel-Schneider, and D. Nardi, *The description logic handbook: Theory, implementation and applications.* Cambridge university press, 2003.

[8] A. Horn, "On sentences which are true of direct unions of algebras," *The Journal of Symbolic Logic*, vol. 16, no. 1, pp. 14–21, 1951. DOI: 10.2307/2268661.

[9] D. M. Gabbay, I. Hodkinson, and M. Reynolds, *Temporal Logic (Vol. 1): Mathematical Foundations and Computational Aspects.* New York, NY, USA: Oxford University Press, Inc., 1994, ISBN: 0-19-853769-7.

[10] C. Baral, T. Eiter, M. Bjäreland, and M. Nakamura, "Maintenance goals of agents in a dynamic environment: Formulation and policy construction," *Artif. Intell.*, vol. 172, no. 12-13, pp. 1429–1469, 2008. DOI: 10.1016/j.artint.2008.03.005.

[11] Y. Zhang and Y. Ding, "CTL model update for system modifications," *Journal of Artificial Intelligence Research*, vol. 31, pp. 113–155, 2008. DOI: 10.1613/jair.2420.

[12] M. J. Fischer and R. E. Ladner, "Propositional dynamic logic of regular programs," *Journal of computer and system sciences*, vol. 18, no. 2, pp. 194–211, 1979. DOI: 10.1016/0022-0000(79)90046-1.

[13] M. M. Ribeiro, R. Wassermann, G. Flouris, and G. Antoniou, "Minimal change: Relevance and recovery revisited," *Artificial Intelligence*, vol. 201, pp. 59–80, 2013. DOI: 10.1016/j.artint.2013.06.001.

[14] P. T. Guerra and R. Wassermann, "On the uncomputability of partial meet contraction for linear-time temporal logic," *South American Journal of Logic*, vol. 4, pp. 385–406, 2018.

[15] J. S. Ribeiro, A. Nayak, and R. Wassermann, "Towards belief contraction without compactness," in *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018*, 2018, pp. 287–296.

[16] J. S. Ribeiro, A. Nayak, and R. Wassermann, "Belief change and non-monotonic reasoning sans compactness," in *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019*, 2019, pp. 3019–3026. DOI: 10.1609/aaai.v33i01.33013019.

[17] J. S. Ribeiro, A. Nayak, and R. Wassermann, "Belief update without compactness in non-finitary languages," in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019*, 2019, pp. 1858–1864. DOI: 10.24963/ijcai.2019/257.

[18] G. Flouris, "On belief change and ontology evolution," Ph.D. dissertation, University of Crete, 2006.

[19] M. M. Ribeiro, *Belief revision in non-classical logics*. Springer Science & Business Media, 2013. DOI: 10.1007/978-1-4471-4186-0.

[20] D. Lewis, *Counterfactuals*. Oxford, 1976.

[21] P. Gärdenfors and D. Makinson, "Nonmonotonic inference based on expectations," *Artificial Intelligence*, vol. 65, no. 2, pp. 197–245, 1994. DOI: 10.1016/0004-3702(94)90017-5.

[22] T. A. Meyer, W. A. Labuschagne, and J. Heidema, "Refined epistemic entrenchment," *Journal of Logic, Language and Information*, vol. 9, no. 2, pp. 237–259, 2000. DOI: 10.1023/A:1008316802976.

[23] M. L. Ginsberg, "Counterfactuals," *Artif. Intell.*, vol. 30, no. 1, pp. 35–79, 1986. DOI: 10.1016/0004-3702(86)90067-6.

[24] D. J. Lehmann and M. Magidor, "What does a conditional knowledge base entail?" *Artif. Intell.*, vol. 55, no. 1, pp. 1–60, 1992. DOI: 10.1016/0004-3702(92)90041-U.

[25] A. Grove, "Two modellings for theory change," *Journal of philosophical logic*, vol. 17, no. 2, pp. 157–170, 1988. DOI: 10.1007/BF00247909.

[26] S. Lindström and W. Rabinowicz, "Epistemic entrenchment with incomparabilities and relational belief revision," in *The logic of theory change*, Springer, 1991, pp. 93–126. DOI: 10.1007/BFb0018418.

[27] D. Makinson and P. Gärdenfors, "Relations between the logic of theory change and nonmonotonic logic," in *The logic of theory change*, Springer, 1991, pp. 183–205. DOI: 10.1007/BFb0018421.

[28] P. Peppas, "Belief revision," in *Handbook of Knowledge Representation*, 2008, pp. 317–359. DOI: 10.1016/S1574-6526(07)03008-8.

[29] D. Zhang and N. Y. Foo, "Infinitary belief revision," *J. Philosophical Logic*, vol. 30, no. 6, pp. 525–570, 2001. DOI: 10.1023/A:1013356315540.