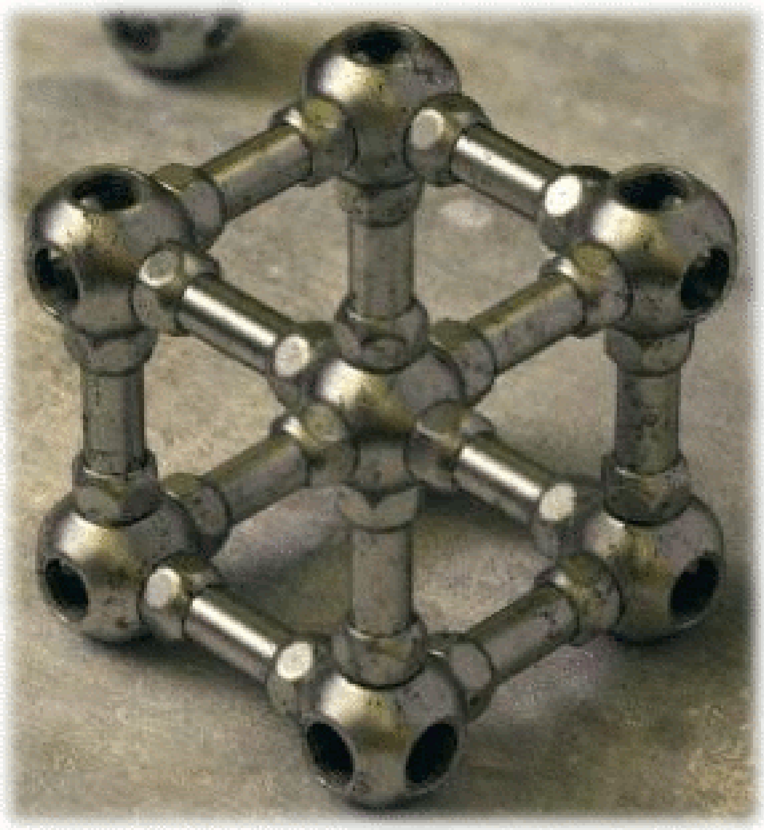

On AGM for Non-Classical Logics [1]

Author(s):
Renata Wassermann

[1]This work was supported by Fapesp Project LogProb, grant 2008/03995-5, São Paulo, Brazil.

On AGM for Non-Classical Logics

Renata Wassermann (renata@ime.usp.br)
Department of Computer Science
Institute of Mathematics and Statistics
University of São Paulo, Brazil

Abstract. The AGM theory of belief revision provides a formal framework to represent the dynamics of epistemic states. In this framework, the beliefs of the agent are usually represented as logical formulas while the change operations are constrained by rationality postulates. In the original proposal, the logic underlying the reasoning was supposed to be supraclassical, among other properties. In this paper, we present some of the existing work in adapting the AGM theory for non-classical logics and discuss their interconnections and what is still missing for each approach.

Keywords: belief change - partial meet contraction - non-classical logics - AGM

1. Introduction

Belief Revision deals with the problem of modifying a knowledge base in view of new information. A knowledge base can be seen as representing the beliefs of a rational agent.

As an example, suppose an agent believes that Toulouse is in the South of France, that during Summer the South of France is sunny and warm, and that it is Summer now. Then the agent arrives in Toulouse and it is rainy and cold. This new information about the weather contradicts what the agent believed and in order to accommodate it, the agent must give up some of her previous beliefs. The reason why this is not a trivial task is that first, it involves choosing what to give up and this choice usually does not depend purely on logical reasoning, but rather on some kind of preference. In our example, the agent is less likely to give up the belief that Toulouse is in the South of France than the belief that during Summer the South of France is sunny and warm. Second, giving up some beliefs may have indirect consequences: the agent believed that she only needed Summer clothes and will have to give up this belief as well...

In order to be able to formalize belief revision, we need to be able to represent at least four elements:

- *Epistemic states*: what the agent believes.
- *Epistemic attitudes*: which attitudes the agent may have with respect to his beliefs.

- *Input*: how the new information is represented.
- *Change operations*: the ways in which the agent can change his epistemic state.

In the AGM theory [1] , the epistemic state is given by *belief sets*, sets of formulas closed under logical consequence. There are three possible epistemic attitudes: if K is the belief set of the agent, $\alpha \in K$ is interpreted as "the agent believes α"; $\neg\alpha \in K$ is interpreted as "the agent rejects α"; and if neither α nor $\neg\alpha$ are in K, then the attitude towards α is undetermined. The new information is represented as a formula. There are three operations of change: *expansion* consists in adding the new information without worrying about consistency; *contraction* concerns removing a belief from the belief set and *revision* is the addition of new information while preserving consistency.

One of the variations of AGM theory that has received attention in the literature defends the use of *belief bases* instead of belief sets. Belief bases are sets of formulas which are usually not closed under logical consequence. Their use has been defended by several researchers, but with different meanings: while for some, following [36], a belief base is simply a more compact way of representing a belief set given by its closure, the line we follow here regards a base as containing those beliefs with an independent standing, those which are explicitly believed by an agent [20, 28, 46]. The use of belief bases introduces the possibility of representing more epistemic attitudes: given a belief base B, a formula α can be explicitly accepted ($\alpha \in B$), implicitly accepted ($\alpha \in Cn(B)$), explicitly rejected ($\neg\alpha \in B$), implicitly rejected ($\neg\alpha \in Cn(B)$), or undetermined ($\alpha \notin Cn(B)$ and $\neg\alpha \notin Cn(B)$).

Although most people assume that the AGM framework is meant for classical propositional logic, the class of logics to which it can be applied is a bit broader than that. Quoting the original paper [1]:

> "By a consequence operation we mean, as is customary, an operation Cn that takes sets of propositions to sets of propositions, such that three conditions are satisfied, for any sets X and Y of propositions: $X \subseteq Cn(X)$, $Cn(X) = Cn(Cn(X))$, and $Cn(X) \subseteq Cn(Y)$ whenever $X \subseteq Y$. (...) we assume that Cn includes classical tautological implication, is compact (that is, $y \in Cn(X')$ for some finite subset X' of X whenever $y \in Cn(X)$), and satisfies the rule of "introduction of disjunctions in the premises" (that is, $y \in Cn(X \cup \{x_1 \vee x_2\})$ whenever $y \in Cn(X \cup \{x_1\})$ and $y \in Cn(X \cup \{x_2\})$)."

The first three conditions define a *Tarskian consequence operator*. We will refer to the six conditions together as the *AGM assumptions*.

DEFINITION 1. *An inference operator Inf satisfies the AGM assumptions if and only if for every set of formulas X, Y:*

- $X \subseteq \mathit{Inf}(X)$ **(inclusion)**
- $\mathit{Inf}(X) = \mathit{Inf}(\mathit{Inf}(X))$ **(idempotency)**
- $\mathit{Inf}(X) \subseteq \mathit{Inf}(Y)$ *whenever* $X \subseteq Y$ **(monotony)**
- $Cn(X) \subseteq \mathit{Inf}(X)$ **(supraclassicality)**
- $\alpha \in \mathit{Inf}(X')$ *for some finite subset* X' *of* X *whenever* $\alpha \in \mathit{Inf}(X)$ **(compactness)**
- $\alpha \in \mathit{Inf}(X \cup \{\beta_1 \vee \beta_2\})$ *whenever* $\alpha \in \mathit{Inf}(X \cup \{\beta_1\})$ *and* $\alpha \in \mathit{Inf}(X \cup \{\beta_2\})$ **(introduction of disjunction)**

Classical logic clearly satisfies all of the AGM assumptions. Why should one consider other logics for belief revision? One reason is purely theoretical - as a logical exercise it would be interesting to know what happens if we abandon the AGM assumptions on the underlying logic. Which results still hold and which ones do not? But there are many other reasons for leaving the classical realm: we may be interested in logics with better behavior in terms of computational complexity, logics that tolerate some inconsistency, logics that model different kinds of reasoning, logics used in specific domains of Artificial Intelligence, ...

In this paper we give an overview of some of the work that has been done on applying AGM theory to different logics. We will deal with both belief sets and belief bases, making explicit for each logic, whether the results were obtained for belief sets or for bases. The present paper is not intended as an exhaustive survey of the field, instead, we have selected four main collections of systems, grouped by their motivation.

The paper is organized as follows: in the next section, we give the background on AGM theory which is needed. In Section 3, we present generalizations of AGM theory which can be applied to different logics. In the four sections that follow, we present results which were obtained for particular logics: approximate and local reasoning in Section 4, paraconsistent and relevance logics in Section 5, Horn logics in Section 6, and description logics in Section 7. Finally, in Section 8 we discuss the connections between the different approaches and future work.

Notation: Given a language L, we call *inference operation*, denoted by *Inf*, any total function taking sets of formulas to sets of formulas. We reserve Cn to denote the classical propositional consequence operator. The Greek letters $\alpha, \beta, \gamma \ldots$ denote arbitrary formulas; lowercase letters

$p, q, r \ldots$ denote propositional atoms; uppercase letters $A, B \ldots$ denote sets of formulas. For any formula α, $Var(\alpha)$ is the set of propositional letters which occur in α.

2. Background

In this Section, we very briefly review the AGM theory. For more details, the reader is referred to [21, 22, 28].

We start by assuming that the belief state of an agent is represented by a belief set $K = Cn(K)$. We have seen that in the AGM theory, the input is a formula α. The operation of expanding K by α can be defined as:

$$K + \alpha = Cn(K \cup \{\alpha\})$$

The resulting set is a belief set that contains α and that may be inconsistent. The other two change operations, contraction and revision, do not have a unique construction but are constrained by *rationality postulates*. Following the formulation given in [21], the AGM postulates for contraction can be written as:

(K-1) $K - \varphi$ is a belief set (**closure**)
(K-2) $K - \varphi \subseteq K$ (**inclusion**)
(K-3) If $\varphi \notin K$, then $K - \varphi = K$ (**vacuity**)
(K-4) If not $\vdash \varphi$, then $\varphi \notin K - \varphi$ (**success**)
(K-5) $K \subseteq (K - \varphi) + \varphi$ (**recovery**)
(K-6) If $\vdash \varphi \leftrightarrow \psi$, then $K - \varphi = K - \psi$ (**equivalence**)
(K-7) $K - \varphi \cap K - \psi \subseteq K - (\varphi \wedge \psi)$
(K-8) If $\varphi \notin K - \varphi \wedge \psi$, then $K - \varphi \wedge \psi \subseteq K - \varphi$

The first six postulates are considered to be more basic, while the last two deal with contraction by conjunctions. In the rest of the paper, when we refer to the postulates, we mean the six basic postulates, unless explicitly mentioned. The same holds for the revision postulates:

(K*1) $K * \varphi$ is a belief set (**closure**)
(K*2) $\varphi \in K * \varphi$ (**success**)
(K*3) $K * \varphi \subseteq K + \varphi$ (**inclusion**)
(K*4) If $\neg\varphi \notin K$, then $K + \varphi \subseteq K * \varphi$ (**preservation**)
(K*5) $K * \varphi = L$ if and only if $\vdash \neg\varphi$ (**consistency**)
(K*6) If $\vdash \varphi \leftrightarrow \psi$, then $K * \varphi = K * \psi$ (**equivalence**)
(K*7) $K * (\varphi \wedge \psi) \subseteq (K * \varphi) + \psi$
(K*8) If $\neg\psi \notin K * \varphi$, then $(K * \varphi) + \psi \subseteq K * (\varphi \wedge \psi)$

Contraction and revision can be defined in terms of each other through the Levi or the Harper identities. Revision by a belief α corresponds to first contracting by $\neg\alpha$ and then expanding by α:

$K * \alpha = (K - \neg\alpha) + \alpha$ (Levi identity)

In a similar way, contraction by α can be obtained through revision by $\neg\alpha$ followed by the elimination of what was not in the original set:

$K - \alpha = (K * \neg\alpha) \cap K$ (Harper identity)

Several constructions for the change operators have been proposed in the literature. Here we will mostly use the construction that appeared in [1], *partial meet contraction.* This operation is based on selecting some of the maximal subsets of the belief set that do not imply the input formula and taking their intersection. Formally:

DEFINITION 2. *[2] The* **remainder set** $X\bot\alpha$ *of* X *and* α*, where* X *is a set of formulas and* α *a formula, is defined as follows. For any set* Y, $Y \in X\bot\alpha$ *if and only if:*

– $Y \subseteq X$

– $Y \nvdash \alpha$

– *For all* Y' *such that* $Y \subset Y' \subseteq X$, $Y' \vdash \alpha$.

DEFINITION 3. *[1] A* **selection function** *for* X *is a function* γ *such that:*

– *If* $X \perp \alpha \neq \emptyset$*, then* $\emptyset \neq \gamma(X \perp \alpha) \subseteq X \perp \alpha$.

– *Otherwise,* $\gamma(X \perp \alpha) = \{X\}$.

DEFINITION 4. *[1] For any sentence* α*, the operation of* **partial meet contraction** *over a belief set* K *determined by the selection function* γ *is given by:*

$$K -_\gamma \alpha = \bigcap \gamma(K \perp \alpha)$$

THEOREM 1. *[1] Let* $-$ *be a function which, given a formula* α*, takes a belief set* K *into a new belief set* $K - \alpha$*. If the underlying logic satisfies the AGM assumptions, then for every theory* K, $-$ *is a partial meet contraction operation over* K *if and only if* $-$ *satisfies the basic postulates* **(K-1)**-**(K-6)** *for contraction.*

Partial meet revision is defined using the Levi identity and partial meet contraction.

Another operation which had already been defined in [2] can be seen as a special case of partial meet contraction. In a *maxichoice contraction*, the selection function γ selects a single element of $K \perp \alpha$, in case this set is not empty. In [1], it is shown that the following postulate, together with the six basic AGM postulates, fully characterizes maxichoice contractions:

(K-F) If $\beta \in K$ and $\beta \notin K-\alpha$, then $\beta \rightarrow \alpha \in K-\alpha$ (**fullness**)

The following results show the undesirable effects of maxichoice operations:

LEMMA 1. *[2] If $\alpha \in K$ and $K-\alpha$ is defined by means of a maxichoice contraction operation, then for any formula β, either $\alpha \vee \beta \in K-\alpha$ or $\alpha \vee \neg\beta \in K-\alpha$.*

COROLLARY 1. *[2] If a revision operation is defined from a maxichoice contraction by means of the Levi identity, then, for any α such that $\neg\alpha \in K$, $K * \alpha$ will be maximal, i.e., for every formula β, either $\beta \in K * \alpha$ or $\neg\beta \in K * \alpha$.*

Suppose I believe p (that Buenos Aires is the capital of Brazil) and have no idea about q (that the moon is made of cheese). Finding out that $\neg p$ is the case and revising my belief set using a revision based on maxichoice contraction means that I will make a decision as to whether q or $\neg q$.

The other extreme of partial meet operators is *full meet contraction*, also defined in [2], where the whole set $K \perp \alpha$ (if not empty) is selected. Full meet contraction is the only one of the three variations that does not require a selection function.

Full meet contractions can be characterized by the basic postulates together with the following postulate:

(K-I) For all α and β, $K-(\alpha\wedge\beta) = (K-\alpha)\cap(K-\beta)$ (**intersection condition**)

The following results show that full meet contraction deletes beliefs that intuitively should be preserved:

LEMMA 2. *[2] If $\alpha \in K$ and $K-\alpha$ is defined by means of a full meet contraction operation, then $\beta \in K-\alpha$ if and only if $\beta \in K$ and $\beta \in Cn(\neg\alpha)$.*

COROLLARY 2. *[2] If a revision operation is defined from full meet contraction by means of the Levi identity, then, for any α such that $\neg\alpha \in K$, $K * \alpha = Cn(\alpha)$.*

Suppose I believe that p (Buenos Aires is the capital of Brazil) and that q (there is no King of France). When I learn $\neg p$ and revise my belief set using a revision operation based on full meet contraction, I give up the belief that there is no King of France.

When instead of belief sets we use belief bases to represent the agent's epistemic state, we can still use partial meet constructions to obtain contraction and revision operators. However, these operations are not characterized by the same set of postulates as the original AGM constructions. Hansson [25] has provided sets of postulates and representation theorems for partial meet contraction and revision of belief bases.

An alternative construction was proposed in [27], where instead of using the maximal subsets of a base that do not imply a given formula, we rely on the minimal subsets that do imply the input formula. The formal definition of *kernel contraction* is as follows:

DEFINITION 5. *[27] The* **kernel** *operation* $\perp\!\!\!\perp$ *is the operation such that for every set* B *of formulas and every formula* α, $X \in B \perp\!\!\!\perp \alpha$ *if and only if:*

- $X \subseteq B$
- $\alpha \in Cn(X)$
- *for all* Y, *if* $Y \subset X$ *then* $\alpha \notin Cn(Y)$

The elements of $B \perp\!\!\!\perp \alpha$ *are called* α*-kernels.*

DEFINITION 6. *[27] An* **incision function** *for* B *is any function* σ *such that for any formula* α*:*

- $\sigma(B \perp\!\!\!\perp \alpha) \subseteq \bigcup(B \perp\!\!\!\perp \alpha)$, *and*
- *If* $\emptyset \neq X \in B \perp\!\!\!\perp \alpha$, *then* $X \cap \sigma(B \perp\!\!\!\perp \alpha) \neq \emptyset$.

DEFINITION 7. *[27] Let* σ *be an incision function. The* **kernel contraction** *on* B *determined by* σ *is the operation* $-_\sigma$ *such that for all sentences* α*:*

$B -_\sigma \alpha = B \setminus \sigma(B \perp\!\!\!\perp \alpha)$

Hansson has provided in [27] a set of postulates and a representation theorem for this operation and has shown that the operation is more general than partial meet base contraction.

3. General results

Some of the results obtained for particular logics proved to hold for a larger class of logics. As an example, although the operations proposed in [29] were originally coined for local reasoning (see Section 4), the corresponding representation theorems hold for several different logics. To make this clear, we reproduce below one of them:

THEOREM 2. *[29] Let Inf satisfy monotony and compactness. Then $-$ is an operator of partial meet contraction on B based on Inf if and only if for all sentences α:*

- *If $\alpha \notin Inf(\emptyset)$, then $\alpha \notin Inf(B - \alpha)$ (**success**)*
- *$B - \alpha \subseteq B$ (**inclusion**)*
- *If $\beta \in B \setminus (B - \alpha)$, then there is some B' such that $B - \alpha \subseteq B' \subseteq B$, $\alpha \notin Inf(B')$ and $\alpha \in Inf(B' \cup \{\beta\})$ (**relevance**)*
- *If for all subsets B' of B, $\alpha \in Inf(B')$ if and only if $\beta \in Inf(B')$, then $B - \alpha = B - \beta$ (**uniformity**)*

For revision, one more property is needed, which is a property of a pair *Inf*,α and states that according to the inference operator *Inf*, the formula α cannot be responsible for the derivation of $\neg\alpha$:

If $\neg\alpha \in Inf(B \cup \{\alpha\})$, then $\neg\alpha \in Inf(B)$.**(α-local non-contravention)**

The representation theorems are basically the same as in the classical case, proved in [25], but with the classical consequence operator Cn replaced with a generic operator satisfying the properties specified in the theorem. What is important is that the same consequence operator is used both in the construction, to decide whether a subset of the belief base implies or not a given formula, and in the postulates. Thus, the meaning of the postulates changes according to the logic used, as discussed in [50]. For example: the **success** postulate for contraction says that after a contraction by α, the resulting belief base should not imply α. It may well be that the resulting base classically implies α, but not in the particular logic used.

For belief sets, there are some interesting general results to be extracted from work in Description Logics and Horn Logics. Flouris [17] has characterized exactly the logics that admit an AGM contraction operator (one that satisfies the six AGM postulates) through the property of *decomposability*:

For all sets of formulas A, B, such that $Cn(\emptyset) \subset Cn(B) \subset Cn(A)$, there exists a set of formulas C such that $Cn(C) \subset Cn(A)$ and $Cn(A) = Cn(B \cup C)$.

Flouris, Plexousakis and Antoniou have shown that the problem lies in the **recovery** postulate:

THEOREM 3. *[18] Every Tarskian logic admits a contraction operator that satisfies the AGM postulates without* **recovery**.

Based on this result, it was shown in [41] that by substituting the **relevance** postulate for **recovery**, one can obtain contraction operators for more logics, even those that are not decomposable. An interesting property is that the two sets of postulates (the original AGM and AGM with **relevance** instead of **recovery**) are equivalent for the logics satisfying the AGM assumptions, which makes the latter a real extension.

THEOREM 4. *[41] Every Tarskian logic admits a contraction operator that satisfies the AGM postulates with* **relevance** *instead of* **recovery**.

Finally, **relevance** can be used to characterize partial meet contraction:

THEOREM 5. *[41] For every belief set K closed under a Tarskian consequence operator, $-$ is a partial meet contraction operation over K if and only if $-$ satisfies the postulates* **(K-1)**-**(K-4)**, **(K-6)** *and* **(relevance)**.

Most papers on belief change, when it comes to construction, concentrate on contraction, since it is supposed to be a more fundamental operation and then rely on the Levi identity for constructing revision. The Levi identity states that revision by α can be obtained through contraction by $\neg\alpha$ followed by expansion. This works well for classical logic, but several non-classical logics lack full classical negation. In many description logics there is no way to express the negation of an arbitrary concept inclusion. In Horn logic, also a fragment of classical logic, with classical semantics, the negation of an arbitrary formula very often falls outside the language. In other cases, such as paraconsistent logics and approximate logics, the semantic of negation is not classical, so it is not always clear whether the Levi identity makes sense at all. There were attempts both in description logics [43, 42] and in Horn logics [14] to provide a construction for revision that does not depend on negation. For belief sets, none of them is general enough (see Section 7).

For belief bases, six different constructions were studied in [43], each with an associated set of postulates that fully characterized it. The main idea is to use a version of the *reversed Levi identity* [26] where instead of first contracting by the negation and then expanding, we first expand the belief base and then contract by the contradiction. In this way, we do not need an explicit negation of the input formula. The difference between the six constructions is whether they are based on kernel or partial meet contraction and the way in which they assure (or not) that the input is an element of the resulting set. The results can be applied for any compact and monotonic logic.

In the next sections, we will recall some work on adapting the AGM theory to other logics. Although we have tried to divide the approaches mentioned into four main groups, the borders between the groups are not always clear. We will discuss this in Section 8.

4. Local and Approximate Reasoning

One of the main criticisms against the use of logic for formalizing the reasoning of an agent is the computational complexity. Even propositional logic is intractable. However, the use of logic as a representation language is much less controversial. In this section we will present the use of approximations to deal with the complexity of reasoning and the applications to belief revision.

The approaches we will see here are based on the idea of local reasoning. The main intuition is that an agent with bounded resources is not able to take everything he knows simultaneously into account. Usually he does not even need to, in order to perform an inference. Each logic provides a different method for selecting which part of the knowledge base should be taken into account. In this section, we only consider operations on belief bases.

4.1. Local compartments

In [29], a notion of relevant compartments around a given formula was introduced. This notion tries to capture a kind of logical relevance: the compartment of a belief base around a formula α contains the formulas of the belief base that are involved in proving or refuting α. Compartmentalization is achieved by taking the union of all the minimal subsets of the base that imply α and those that imply $\neg\alpha$ and removing the inconsistent ones (since they imply everything). For the formalization of the logical compartments, we use the notation $B \perp\!\!\!\perp \alpha$ introduced in Definition 5:

DEFINITION 8. *[29] The function c is the compartmentalization function based on Cn if and only if, for all sets of formulas A, B:*

$c(A, B) = \bigcup_{\alpha \in A} c(\alpha, B)$,

$$\textit{where } c(\alpha, B) = \begin{cases} \emptyset \textit{ if } \alpha \in Cn(\emptyset) \textit{ or } \neg\alpha \in Cn(\emptyset) \\ \bigcup(((B \perp\!\!\!\perp \alpha) \cup (B \perp\!\!\!\perp \neg\alpha)) \setminus (B \perp\!\!\!\perp \bot)) \textit{ otherwise.} \end{cases}$$

Using compartments, a consequence operator Cn_α is defined as $Cn_\alpha(X) = Cn(c(\alpha, X))$, where $c(\alpha, X)$ is the compartment of X around α. This consequence operator is used instead of the classical one to define operations of contraction and revision of belief bases. For partial meet operators, if one wants to contract a belief base B by α, the remainder sets to be considered are those maximal subsets of B that do not imply α according to Cn_α, i.e., the maximal subsets X of B such that $\alpha \notin Cn(c(\alpha, X))$. The representation theorems provided in this context are very general, as mentioned in Section 3 and can be used with any compact and monotonic consequence operator that satisfies α-local non-contravention as is the case of the local consequence operator defined in terms of compartments.

A problem with the compartments approach is that computing the compartments is at least as hard as performing the belief change operations. Hence, although the notion of compartment and local change seem more adequate to model a realistic agent reasoner, in terms of computational complexity, there is no gain. Following the idea of selecting a relevant part of the base in order to apply a change operator, one can think of more efficient methods for achieving this selection. In [49], an abstract relatedness relation $\mathcal{R}$ between formulas is used in order to build a graph that "structures" the belief base. Any relation satisfying reflexivity, symmetry and the property that for any φ, $\mathcal{R}(\varphi, \neg\varphi)$ can be used as a relatedness relation. An example of such relation is atom sharing, i.e., we define $\mathcal{R}(\alpha, \beta)$ if and only if the two formulas have a propositional variable in common ($Var(\alpha) \cap Var(\beta) \neq \emptyset$).

DEFINITION 9. *[49] Let B be a belief base and $\mathcal{R}$ be a relation between formulas. An $\mathcal{R}$-**path** between two formulas φ and ψ in a belief base B is a sequence $P = (\varphi_0, \varphi_1, ..., \varphi_n)$ of formulas such that:*

- $\varphi_0 = \varphi$ *and* $\varphi_n = \psi$
- $\{\varphi_1, ..., \varphi_{n-1}\} \subseteq B$
- $\mathcal{R}(\varphi_i, \varphi_{i+1})$, $0 \leq i < n$.

The **length** *of a path* $P = (\varphi_0, \varphi_1, ..., \varphi_n)$ *is* $l(P) = n$

The idea of having a relatedness relation with respect to a given belief base as opposed to the whole language was first proposed by Rodrigues in [45].

The local inference operator takes a natural number i as a parameter and performs classical inference considering only those formulas from the belief base that are at distance at most i from the input, where the distance is given by the length of the shortest path between two formulas. This generates an inference operator that can be used together with the results in [29]:

$C^i_\alpha(B) = Cn(X)$,
where $X = \{\beta \in B | \beta \text{ is at most at distance } i \text{ from } \alpha\}$

The notion of approximation becomes apparent here: for each i, C^i_α is sound but possibly incomplete with respect to Cn and as we increase i, we get closer to classical inference.

Makinson has recently presented an overview of different notions of relevance and how they relate to AGM revision [34]. He classified the notions according to their syntax and language dependencies. These different notions can be used as the relatedness relation $\mathcal{R}$ in order to form the relevant compartments.

4.2. Belief Sequences

A related approach was proposed in [9]. Chopra, Georgatos and Parikh define a notion of relevance between two formulas which is less syntactic than simple atom sharing: they consider atom sharing in the *minimal language* needed to express the content of the two formulas. If $Var(\alpha)$ is the set of propositional letters of a formula α, $Var_{min}(\alpha)$ is defined as the minimal set of propositional letters that occur in a formula equivalent to α:

$$Var_{min}(\alpha) = min\{Var(\beta) | \vdash \beta \leftrightarrow \alpha\},$$

where min is defined with respect to set inclusion. Note that the minimum is unique provided that either $\top$ or $\bot$ are in the language.

Thus, although $(p \vee \neg p) \wedge q$ and $\neg q \rightarrow p \wedge q \rightarrow p$ share the atoms p and q, they are not considered related since the first is equivalent to q and the second to p. This relation could be used as the relatedness relation $\mathcal{R}$ in the relevance graph approach.

The main difference between this work and the previous two is that in this one, revision is not defined in the AGM sense. The epistemic states are represented by sequences of formulas, instead of sets. Revision by α is the simple concatenation of α at the end of the sequence. The beliefs of the agent are given not by the classical closure of the set of beliefs, but buy a prioritized inference on the sequence, that takes into account the linear ordering (more recent beliefs have higher priority) and relevance. This can be seen as an example of the *vertical mode* of revision as opposed to the AGM *horizontal mode* discussed in [46].

The inference is achieved by reordering the sequence according to relevance to α (a formula is more relevant to α than another if it is closer to α in terms of the path through relevance links) using the original temporal ordering to break ties. Then a maximal consistent subset of formulas is computed by following the sequence, starting from the most relevant formula and adding to the subset those formulas which are consistent with the ones previously added. The authors show that the revision defined from the relevant sequence using prioritized inference satisfies at least two AGM postulates - **success** and **extensionality**. The others were not checked and should be reformulated for the sequences setting before they can be verified.

4.3. Belief structures

Following the line of local reasoning, Chopra and Parikh [10] state that in a belief change operation only the part of the agent's beliefs which is relevant to the input should be changed. A trivial revision operator, which forgets all the previous beliefs and keeps only the consequences of the input should be avoided, although it satisfies the AGM postulates for revision. Starting from Parikh's notion of splitting language [37], they represent the epistemic state of an agent by a *B-structure*. A B-structure is a tuple of the form $\{(L_1, B_1), ..., (L_n, B_n)\}$ where in each B_i only atoms from L_i occur. To decide whether α is implicitly believed in a B-structure, one looks at the minimal language needed to express α, $Var_{min}(\alpha)$ and sets $B_\alpha = \bigcup\{B_i | L_i \cap Var_{min}(\alpha) \neq \emptyset\}$. If B_α is consistent, and $B_\alpha \vdash \alpha$, then α is implicitly believed. So once more, the inference operator is built by applying classical inference to a subset of the belief base. In this case, the base is divided a priori into sub-bases and the relevant sub-bases are selected for a given input. The authors define two strategies for revision which only change the sub-bases relevant to the input but provide no postulates or representation theorems. The relationship between this approach and the logical compartments approach is partially explored in [38].

4.4. Approximate logics

A very different approach was suggested in [11]. In this work, the approximations of classical logics introduced by Schaerf and Cadoli [47] are applied to operations of contraction and revision of belief bases.

The idea behind the approximations is to generate logics which are weaker than classical propositional logics and for which satisfiability is easier to compute. This is achieved by disconsidering some of the propositional letters involved in a base. The belief based is assumed to be in clausal form. Schaerf and Cadoli propose two different approximations of classical entailment: $\models^1_S$ which is complete but not sound, and $\models^3_S$ which is sound and incomplete. Both take a set S of atoms as a parameter. In an S_1 assignment, if $x \in S$, then $x, \neg x$ are given opposite truth values; if $x \notin S$, then $x, \neg x$ both get the value 0. In an S_3 assignment, if $x \in S$, then $x, \neg x$ get opposite truth values, while if $x \notin S$, $x, \neg x$ do not both get 0, but may both get 1. There are two extreme cases: when S contains all the atoms involved in an entailment check, the approximations coincide with classical entailment; and when $S = \emptyset$, $\models^1_S$ is trivial, i.e., for any α,β, we have $\alpha \models^1_S \beta$ and $\models^3_S$ corresponds to Levesque's logic for explicit beliefs [33]. Schaerf and Cadoli show that testing whether $B \models^3_S \alpha$ or $B \models^1_S \alpha$ takes $O(|B|.|\alpha|.2^{|S|})$ time. Thus, the idea of approximations is to start with a small S. If $B \models^3_S \alpha$, then we already know that $B \models \alpha$, since S_3 approximation is sound. If $B \not\models^1_S \alpha$, then $B \not\models \alpha$, since S_1 is complete. In case neither $B \models^3_S \alpha$ nor $B \not\models^1_S \alpha$ for a given S, we have to add atoms to S.

These two entailments can be used instead of classical logic in order to provide constructions and rationality postulates for belief base operations. In [11], it is shown that the approximate entailments satisfy the properties needed for the general results in [29] to hold. If we look at contraction, an operation defined using S_1 discards more formulas than one using classical entailment while the same operation using S_3 discards less. This means that we can use them as lower and upper bounds to what the classical result should be. If we have more resources available, we can increase the size of S and obtain a finer interval.

5. Paraconsistent and Relevant Logics

A logic is said to be *paraconsistent* if it is not the case that for any formulas α and β it holds that $\alpha, \neg\alpha \vdash \beta$. *Relevant logics* (or relevance logics) do not have such a clear cut definition, but in general, these logics reject implications where the antecedent and the consequent are not related. One of the rejected implications is $\vdash (p \wedge \neg p) \rightarrow q$, which means that relevant logics are also paraconsistent. In this section, we present some of the existing approaches for adapting AGM to paraconsistent and relevant logics. It is interesting to note that the inference operators introduced in Section 4 are all paraconsistent and relevant. The reason for them to be in a separate section is that they have a different motivation. The works presented in this section deal with belief sets, although usually not closed under classical consequence.

The general motivation for adopting a paraconsistent logic varies: while some authors claim that inconsistencies exist in the world, other appeal to the fact that we do not always know how to solve an inconsistency, so we should learn how to reason in their presence. In classical logic, inconsistency leads to trivialization, i.e., in the presence of inconsistency one can derive any formula of the language. This means that all information is lost. In a paraconsistent logic, inconsistency does not directly lead to trivialization.

5.1. The C hierarchy

Da Costa and Bueno [12] argue that most of the belief revision literature following Levi and Gärdenfors simply take for granted that a belief state must be consistent. This is due to the fact that in classical logic, inconsistency leads to trivialization and so means lack of information. Gärdenfors considers an inconsistent belief state as an "epistemic hell" to be avoided by all means [21].

Da Costa and Bueno do not defend the existence of "true contradictions" but motivate the use of paraconsistent logics by the need to account for reasoning in the presence of inconsistency, and avoid trivialization. Then a question that arises is whether there is need for revision, or whether one should just keep all the information in the belief state. The logics they use (the C hierarchy of da Costa [13]) all allow for inconsistency, but can be trivialized by some contradictions. These are the ones to be avoided in revision. The logics in the C hierarchy avoid trivialization by changing the behavior of negation.

As an example of such logics, we look at Da Costa's system C_1, the first logic in the hierarchy. Some formulas are said to be *well behaved* and for them, negation behaves classically. For formulas which are not well behaved, it may happen that both the formula and its negation are assigned the value true. Thus, having α and $\neg\alpha$ in a belief set does not necessarily lead to trivialization. But having α and $\neg\alpha$ together with $\neg(\alpha \wedge \neg\alpha)$, which is not a tautology in C_1, does mean that one can infer any formula of the language. As one moves higher in the hierarchy, the logics become weaker and thus, harder to trivialize.

Da Costa and Bueno suggest a slight change in he AGM postulates for contraction and revision, where the consequence operation is one of the C_i in the hierarchy and any mention of inconsistency is substituted by trivialization. What is interesting in their proposal is that when one comes across a trivial belief set, there are at least two alternatives:

change the underlying logic for one higher in the hierarchy or revise. The authors do not provide any construction for the operations.

5.2. Four-valued logics

In their paper [40] Slaney and Restall defend the use of first degree entailment [4] instead of classical logic for closing belief sets. The version of logic they use is based on the idea of allowing for subsets of $\{T, F\}$ as truth values, so that formulas can be true, false, both or neither true nor false. The language is formed by atoms and the connectives $\wedge, \vee, \neg$. Valuations on the atoms are extended to formulas in the following way:

$T \in v(\alpha \wedge \beta)$ iff $T \in v(\alpha)$ and $T \in v(\beta)$

$F \in v(\alpha \wedge \beta)$ iff $F \in v(\alpha)$ or $F \in v(\beta)$

$T \in v(\alpha \vee \beta)$ iff $T \in v(\alpha)$ or $T \in v(\beta)$

$F \in v(\alpha \vee \beta)$ iff $F \in v(\alpha)$ and $F \in v(\beta)$

$T \in v(\neg\alpha)$ iff $F \in v(\alpha)$

$F \in v(\neg\alpha)$ iff $T \in v(\alpha)$

Finally, $\alpha \in Cn_{\text{fde}}(X)$ if and only if any valuation making each element of X at least true makes α at least true.

Restall and Slaney first look at the AGM postulates, and abandon **recovery** for contraction. The axioms for revision remain the same. They are obtained from the ones for contraction via the Levi identity, as in the classical case. The Harper identity, however, does not always produce a contraction. The authors examine three classical constructions for contraction operators and their relation to the postulates in the setting of first-degree entailment. First they show that the classical representation result for a construction based on entrenchment [21] goes through even without **recovery**. For partial meet constructions they need an adaptation, dropping the maximality of remainders, then, according to them, the representation theorem holds. The consequence of abandoning the maximality of remainders is that any set contained in a maximal remainder could be the result of contraction. There is no minimal change involved, thus **recovery** is not replaced by any other postulate. But then, as was done in another context by [6] (see section 6) if we allow for non-minimal remainders, there is no need to take their meet, one can simply set one of the remainders as the result of contraction. The last construction they examine is based on systems of spheres [23]. They propose a modification of the definition of spheres so that they are not based on possible worlds, which correspond to classical valuations, but on *prime theories*, which correspond to the four-valued setting described above. For this construction there is no representation result.

The authors end the paper by suggesting a way in which inconsistencies can be solved gradually and locally. This can be achieved by fixing a vocabulary and requiring that the beliefs built from this vocabulary are the most deeply entrenched, or hardest to be given up. The belief set restricted to this vocabulary should be consistent.

Independently and with a different motivation, Lakemeyer and Lang also proposed the use of four-valued logics for AGM revision. In [32], they present postulates very similar to the ones presented by Restall and Slaney, with the main concern being computational tractability instead of verifying representation theorems.

5.3. Rejection sets

Mares [35] proposes to exchange the notion of consistency for what he calls coherence. He defines the belief state of an agent to consist of two sets, one for accepted beliefs and one for rejected beliefs. The coherence condition states that the two must not overlap. He uses the relevant logic R [5] as a paraconsistent logic, but all the results are valid for other relevant logics between B and $RM3$. Acceptance sets are closed under logical consequences. Rejection sets are closed under *downward consequences*, i.e., a rejection set contains every formula that entails any of its formulas. Mares proceeds to define constructions for contraction, expansion, package contraction and revision. While the constructions for the acceptance sets follow the AGM construction of partial meet, for rejection sets things are redefined in a dual way, using downward consequences.

Given a rejection set Δ, the rejection remainders of Δ and α, denoted by $\Delta\top\alpha$ are the maximal subsets of Δ such that none of their elements are implied by α. Rejection contraction is defined by selecting some of the rejection remainder and taking their intersection. Then revision is defined in two different ways: revising by a new acceptance and revising by a new rejection. Since acceptance and rejection sets are not allowed to overlap in a coherent state, accepting a new belief usually involves contracting the rejection set and vice-versa. The revisions satisfy four of the six basic AGM postulates for revision and fail **preservation** and **consistency**, as expected, since the underlying logic is paraconsistent.

5.4. Other proposals

Priest [39] suggests a model for belief change that is supposed to be more general than AGM. He proposes to take for each set of beliefs and a set of criteria, a measure of how good the set is according to each criterion. Then the different measures must be amalgamated according to weights. A partial ordering is given by this amalgamated measures. The model fails every single AGM postulate for revision and Priest argues that it should be so. Actually, he is concerned with an operation quite different from AGM revision, since he wants to account for "conceptual breakthroughs". This is a very general paper and does not adopt any particular logic.

Tanaka [48] proposes to adopt paraconsistent logics and allow for inconsistent belief sets. He proposes alternative Grove like sphere systems, that as in [40] uses prime theories instead of possible worlds. In the paper he investigates what happens when classical logics are substituted by four different systems: the relevant logic B [5], Da Costa's system C_ω [13], a variation of it called C_n and a non-adjunctive logic [30]. He looks at what happens concerning his version of Grove systems and the AGM postulates for contraction and revision in each case. Considering the eight AGM axioms for contraction, they all hold for C_n, while for the logics B, C_ω and non-adjunctive, all contraction postulates hold except for maybe **recovery**, which is left uncertain, i.e., the paper neither proves that it holds nor that it does not. For the revision postulates, **consistency** does not hold in any of the four systems and for non-adjunctive logics, K*7 does not hold and K*8 is not shown to hold or not.

6. Horn Logics

A fragment of classical logics that deserves attention in the Artificial Intelligence community is Horn logics, due to its good computational properties. A *Horn clause* is a disjunction of literals where at most one is positive. A *Horn formula* is a conjunction of Horn clauses. Horn logic inherits the semantics of classical propositional logic and all the connectives are interpreted in the classical way. The Horn consequence operator Cn^h produces those consequences of a set that are expressible in Horn logics, i.e., $Cn^h(X) = \{\alpha \in Cn(X) | \alpha \text{ is a Horn formula}\}$.

In [14], Delgrande suggests two different ways in which a remainder set can be defined, one based on entailment and the other on inconsistencies. Entailment based remainder sets, or *e-remainders*, are defined in the usual way, as maximal subsets of a belief set (but here a belief set is closed under Horn consequence, instead of classical) not implying a given formula. Inconsistency based remainder sets, or *i-remainders* are defined as maximal subsets of a belief set which are consistent with a given formula. In principle, partial meet constructions based on i-remainders should be used in order to define revision while e-remainder constructions are suitable for partial meet contraction. Delgrande gives a set of rationality postulates that the operations satisfy, but the postulates are not strong enough to characterize them. Basically, the operations satisfy five of the six basic AGM postulates for contraction, the exception being **recovery**. However, the operations have some undesirable properties. One of them is that for any p not mentioned in K, we have $(K - \phi) + p \vdash \phi$. This holds for any operation built on e- or i-remainder sets.

In [6], Booth, Meyer, and Varzinczak show that the construction based on e-remainders lacks a property that partial meet constructions for classical logics have, which they called *convexity*. In classical logic, for a given belief set K and a formula α, any set that contains the full meet contraction of K by α and is contained in one of the possible outcomes of maxichoice contraction (i.e., any set K' such that $\bigcap(K \bot \alpha) \subseteq K' \subseteq X \in K \bot \alpha$) can be obtained by a partial meet construction. In Horn logics, as the following example shows, convexity does not hold:

EXAMPLE 1. **[6]** *Let* $K = Cn\{\neg p \vee q, \neg q \vee r\}$. *It is easy to verify that, for the e-contraction of* $\neg p \vee r$, *maxichoice yields either* $K^1_{mc} = Cn^h\{\neg p \vee q\}$ *or* $K^2_{mc} = Cn^h\{\neg q \vee r, \neg p \vee \neg r \vee q\}$, *that full meet yields* $K_{fm} = Cn^h\{\neg p \vee \neg r \vee q\}$, *and that these are the only three possibilities for partial meet e-contractions. Now consider the Horn belief set* $K' = Cn^h\{\neg p \vee \neg q \vee r, \neg p \vee \neg r \vee q\}$. *It is clear that* $K_{fm} \subseteq K' \subseteq K^2_{mc}$, *but there is no partial meet e-contraction yielding* K'.

Booth et al. then propose the use of *infra-remainder sets*, defining those as any set between full meet and maxichoice contractions. The outcome of contraction is defined as one infra-remainder, selected by some function. In [7], a representation result is given and it is shown that this operation is equivalent to performing a kernel contraction and closing the result under Horn consequence. Although this operation satisfies convexity, it also suffers from the same drawback as Delgrande's operations: adding any atom not mentioned in K will bring back the contracted formula.

It must be noted that these problems only arise when considering belief sets. On Horn belief bases, both e-contraction and infra-remainder contraction work perfectly well, as defined in [15] and [7]. For belief sets, another construction is proposed in [15], based on weakening the notion of e-remainder set.

DEFINITION 10. *[15] Let H be a Horn belief set, and let ϕ be a Horn formula. We define $H \Downarrow_e \phi$ as: $H' \in H \Downarrow_e \phi$ iff $H' = H \cap m$, where m is some maximal Horn theory that does not contain ϕ.*

We call $H' \in H \Downarrow_e \phi$ a weak remainder set *of H and ϕ.*

EXAMPLE 2. *[15] Consider a language containing only three atoms, a, b,and c, let $H = Cn^h(a \wedge b)$, $\phi = a \wedge b$. For $m = Cn^h(a \wedge \neg b \wedge c)$, we have that $H \cap m = Cn^h(a)$, since H and m are both closed under Horn consequence. Note that using classical logic, $K = Cn(a \wedge b)$ and $m = Cn(a \wedge \neg b \wedge c)$, gives $Cn(K \cap m) = Cn(a \wedge (b \vee c))$.*

In classical logics, each remainder set in $K \perp \alpha$ corresponds, semantically to the models of K together with one single counter-model of α. This is not true for e-remainders. The idea of weak-remainders is to restore this correspondence. But one cannot simply add a counter-model, or the resulting theory may not be expressible in Horn logic. If one identifies a model with the set of atoms true in it, the models of a Horn theory are closed under intersection. A weak remainder set is equivalent to adding a counter-model of α to the models of H and closing them under intersection. Representation results are given both for maxichoice and for partial meet constructions using weak remainders instead of e-remainders. It is interesting to note that, as shown in [15], weak remainders and infra-remainders are incomparable, neither being more general than the other.

In [3], the authors propose to use a rule base language as the underlying logic for revision and contraction. The language L_W contains two kinds of formulas: literals and rules of the form $a_1 \wedge a_2 \wedge ... \wedge a_n \rightarrow b$ where the a_i and b are literals. This is slightly more general than Horn, as we may have $\neg p \rightarrow q$ as a formula. However, there is only one derivation rule, a form of generalized modus ponens:

From $a_1, a_2, ..., a_n$ and $a_1 \wedge a_2 \wedge ... \wedge a_n \rightarrow b$ derive b.

The consequence operator C_W based on this derivation is weaker than taking the classical consequences which are in the language, i.e., in general $C_W(X) \neq Cn(X) \cap L_W$. This can be seen by taking $p \rightarrow q$ and $\neg p \rightarrow q$, which in classical logic imply q but not according to the operator C_W. Alechina et al. define a contraction algorithm based on Truth Maintenance Systems and show that it satisfies the AGM basic postulates for contraction without **recovery**. As revision cannot be defined using the Levi identity, the authors use a semi-revision (expansion followed by contraction of the contradiction). This operation satisfies the basic AGM postulates for revision without **success**.

7. Description Logics

Description logics (DL) belong to a family of logics used to represent terminological knowledge. A knowledge base in description logics is a combination of two distinct sets: a TBox for terminological knowledge (knowledge about concepts) and an ABox for assertional knowledge (knowledge about individuals). Starting from a set of atomic concepts and roles, new concepts and roles can be formed using the constructors in the language. The set of constructors and axiom types allowed in a particular description logic is what distinguishes it from the others.

The logic $\mathcal{ALC}$, for example, given concepts C and D and role R, allows for union ($C \sqcup D$), intersection ($C \sqcap D$), complement ($\neg C$) and value restrictions ($\forall R.C$ and $\exists R.C$). Besides the constructors, $\mathcal{ALC}$ allows for concept subsumption ($C \sqsubseteq D$), concept equivalence ($C \equiv D$), concept assertions ($C(x)$) and role assertions ($R(x, y)$) as axioms. Description logics are equivalent to tractable fragments of first-order logic, where a concept can be seen as a monadic predicate and a role as a binary predicate.

There have been several recent attempts to apply belief revision to description logics. One of the first was [17], where belief set and belief base approaches were examined. In his thesis, Flouris shows that AGM is not directly applicable to several description logics, only to those that satisfy the property of decomposability (see Section 3). Logics which are not decomposable do not admit a contraction operator that satisfies the six basic AGM postulates. Flouris shows that several important DLs do not satisfy this property, not even $\mathcal{ALC}$ with a non-empty ABox. An example of a logic that does satisfy the property is $\mathcal{ALC}$ with an empty ABox, provided there are infinite roles.

In order to check whether a given description logic admits an AGM contraction operator, the following result can be used:

THEOREM 6. *[19] Any description logic which admits:*

- *At least two role names and one concept name*
- *At least one of the operators $\forall$, $\exists$, $(\geq_n)$, $(\leq_n)$ for some n*
- *Any (or none) of the operators $\neg$, $\sqcup$, $\sqcap$, $^-$, $\bot$, $\top$, $\{...\}$*
- *Only the connective $\sqsubseteq$ applicable to both concept and roles*

does not admit an AGM contraction operator.

As the theorem above shows, most expressive description logics are not AGM-compliant.

Flouris then proposes to substitute **recovery** by an alternative postulate:

If $(K - \alpha) + \alpha \subset Cn(Y \cup \{\alpha\})$ for some $Y \subseteq K$, then $Cn(\emptyset) \subset Cn(\{\alpha\}) \subseteq Cn(Y)$.

The intuition behind this postulate is that instead of requiring that $(K - \alpha) + \alpha$ is equal to K, the resulting set is only required to be maximal (thus preserving as much information as possible), in the sense that, if there was some subset Y of K that when expanded by α would give a "larger" set than $(K - \alpha)$, the closure of this Y would necessarily contain α and hence not be suitable as a result of contraction by α.

Flouris, Plexousakis and Antoniou show that for logics that admit an AGM contraction operator, the set of postulates obtained by substituting **recovery** by their new postulate is equivalent to the original set.

In [41], it was noted that instead of the above postulate, **relevance**, which appears in Theorem 2, could be used to replace **recovery**. Together with the other five postulates, **relevance** fully characterizes partial meet contraction. This result can be used with most description logics such as $\mathcal{ALC}$ and more expressive description logics such as $\mathcal{SHIF}(\mathcal{D})$ and $\mathcal{SHOIN}(\mathcal{D})$, which are the underlying logics of OWL-Lite and OWL-DL.

For revision things get more complex, since we cannot simply rely on the Levi identity and construct revision from contraction. The Levi identity requires contracting by the negation of the input formula, but most description logics do not admit the negation of an arbitrary formula. This means that revision must be constructed directly, without relying on contraction.

One such result, following the line of [29] in making precise statements about the conditions on the logic was obtained for belief sets:

DEFINITION 11 (Maximally consistent set w.r.t α). *[14] If α is consistent then $X \in K \downarrow \alpha$ iff:*

- $X \subseteq K$
- $X \cup \{\alpha\}$ *is consistent*
- *For all X', if $X \subset X' \subseteq K$ then $X' \cup \{\alpha\}$ is inconsistent*

If α is inconsistent then $K \downarrow \alpha = \{K\}$.

Consider the following properties of inference operators:

- Whenever K is inconsistent, then for all formulas α, $\alpha \in Cn(K)$ **(inconsistent explosion)**
- For all sets of formulas X, Y and W, $Cn(X \cup (Cn(Y) \cap Cn(W))) = Cn(X \cup Y) \cap Cn(X \cup W)$ **(distributivity)**

THEOREM 7 (Revision without negation). *[42] For any monotonic and compact logic that satisfies inconsistent explosion and distributivity, a revision operator $*$ is a revision without negation $\bigcap \gamma(K \downarrow \alpha) + \alpha$ for some selection function γ if and only if it satisfies* **closure**, **success**, **inclusion**, **consistency**, **relevance** *and* **uniformity**.

The two new properties mentioned in the representation theorem are far less general than monotonicity and compactness. Unfortunately, many description logics do not satisfy distributivity, Horn logic does not either. And clearly, paraconsistent and relevant logics do not satisfy inconsistent explosion. This means that this result is not general enough. In [44] it is shown which description logics satisfy the two properties. A more general characterization of partial meet revision of belief sets in description logics is still needed.

For revision of belief bases, the results in [43] can be applied for most description logics. In [24], an algorithm was proposed and tested for one of the six operations listed in [43], namely, using kernel contraction and not ensuring the **success** postulate. The tests were performed using the logic $\mathcal{SHOIN}$.

Recently, a proposal of revision for the family of logics DL-Lite appeared in [8]. The DL-Lite family underlies one of the profiles for OWL2, and is receiving much attention for being tractable. The paper analyzes several existing operation for the evolution of databases, most of which would be classified as update operators in the sense of [31]. They propose a formula-based operation very close to maxichoice revision, and show that it satisfies three desiderata related to consistency, success and minimal change.

8. Discussion and conclusions

In the previous sections we have seen several different logics for which there are applications or adaptations of the AGM theory. Very often these different approaches are developed for one particular logical system, but can be useful for other logics as well. In this section, we discuss the connections between the different logical systems and what is missing on each side.

As we have already mentioned, all of the logics in Section 4 are paraconsistent. In [16], Schaerf and Cadoli's system S_3 has been extended to full propositional logic and compared to da Costa's C_1. The difference between the two systems is just the set of formulas that behave classically. We can use this fact together with the discussions of the use of C_1 for belief revision in [12] and [48] to improve the results of using S_3 obtained in [11]. Recall that the use of S_3 in [11] is restricted to clausal form and to belief bases, while the two approaches for C_1 deal with belief sets. As suggested in [12], the C hierarchy can be seen as an approximation of classical logic, where for $i < j$, C_i is closer to classical logic than C_j.

Another logic that bears resemblance to S_3 is first-degree entailment. By establishing the correct relation between the two, we can expect that the approaches in [11] and [40] can benefit from each other. The clausal fragment is an interesting one for computational purposes, so clausal belief bases using other paraconsistent logics than S_1 and S_3 should be given attention.

Restall and Slaney [40] have proposed the use of non-maximal remainder sets in order to provide constructions for contraction of belief sets using first-degree entailment. This was also proposed by Booth et al. [6] in the context of Horn logics and later shown to be equivalent to kernel constructions. However, the infra-remainder sets proposed in [6] always contain the result of full meet contraction, while there is no such constraint in the work of Restall and Slaney. It would be interesting to see whether the results can be transferred from one logic to the other, and to compare the non-maximal remainders to the weak remainders of [15].

What most of the non-classical logics in this paper have in common is the fact that negation does not behave classically. For Horn and description logics, given a formula α, there does not necessarily exist a formula in the language which expresses the negation of $\neg\alpha$. This makes the definition of revision an issue, since one cannot rely on the Levi identity. The results for revising belief bases in description logics can probably be used for Horn. However, for belief sets, there is not yet a good proposal.

Acknowledgements

This work has been partially supported by FAPESP under project 08/03995-5 and by CNPq through grant 302514/2007-4. The author would like to thank the anonymous referees for helpful comments.

References

1. Alchourrón, C., P. Gärdenfors, and D. Makinson: 1985, 'On The Logic of Theory Change'. *Journal of Symbolic Logic* **50**(2), 510–530.
2. Alchourrón, C. and D. Makinson: 1982, 'On The Logic of Theory Change: contraction functions and their associated revision functions'. *Theoria* **48**, 14–37.
3. Alechina, N., M. Jago, and B. Logan: 2005, 'Resource-bounded belief revision and contraction'. In: M. Baldoni, U. Endriss, A. Omicini, and P. Torroni (eds.): *Proceedings of the Third International Workshop on Declarative Agent Languages and Technologies (DALT 2005)*, Vol. 31. pp. 118–131.
4. Anderson, a. R. and N. D. Belnap: 1963, 'First degree entailments'. *Mathematische Annalen* **149**(4), 302–319.
5. Anderson, A. R. and N. D. Belnap: 1975, *Entailment: The Logic of Relevance and Necessity, Vol. 1.* Princeton University Press.
6. Booth, R., T. Meyer, and I. J. Varzinczak: 2009, 'New Steps in Propositional Horn Contraction'. In: *Proceedings of the International Joint Conference of Artificial Intelligence (IJCAI).* Pasadena.
7. Booth, R., T. Meyer, I. J. Varzinczak, and R. Wassermann: 2010, 'A Contraction Core for Horn Belief Change: Preliminary Report'. In: *Proceedings of the International Workshop on Non-monotonic Reasoning (NMR).*
8. Calvanese, D., E. Kharlamov, W. Nutt, and D. Zheleznyakov: 2010, 'Evolution of DL-Lite Knowledge Bases'. In: *Proceedings of the 9th International Semantic Web Conference (ISWC 2010).*
9. Chopra, S., K. Georgatos, and R. Parikh: 2001a, 'Relevance Sensitive Nonmonotonic Inference on Belief Sequences'. *Journal of Applied Non-Classical Logics* **11**(1-2).
10. Chopra, S. and R. Parikh: 2000, 'Relevance Sensitive Belief Structures'. *Annals of Mathematics and Artificial Intelligence* **28**(1-4), 259–285.
11. Chopra, S., R. Parikh, and R. Wassermann: 2001b, 'Approximate Belief Revision'. *Logic Journal of the IGPL* **9**(6), 755–768.
12. Costa, N. C. A. D. and O. Bueno: 1998, 'Belief Change and Inconsistency'. *Logique & Analyse* (161-162-163), 31–56.
13. da Costa, N. C. A.: 1963, 'Calculs Propositionnels pour les Systémes Formels Inconsistants'. *Comptes Rendus d'Academie des Sciences de Paris* **257**.
14. Delgrande, J. P.: 2008, 'Horn Clause Belief Change: Contraction Functions'. In: G. Brewka and J. Lang (eds.): *Principles of Knowledge Representation and Reasoning (KR2008).* pp. 156–165.
15. Delgrande, J. P. and R. Wassermann: 2010, 'Horn Clause Contraction Functions: Belief Set and Belief Base Approaches'. In: *Principles of Knowledge Representation and Reasoning (KR2010).*

16. Finger, M. and R. Wassermann: 2004, 'Approximate and Limited Reasoning: Semantics, Proof Theory, Expressivity and Control'. *Journal of Logic And Computation* **14**(2), 179–204.
17. Flouris, G.: 2006, 'On Belief Change and Ontology Evolution'. Ph.D. thesis, University of Crete.
18. Flouris, G., D. Plexousakis, and G. Antoniou: 2004, 'Generalizing the AGM postulates: preliminary results and applications.'. In: *Proceedings of the 10th International Workshop on Non-Monotonic Reasoning (NMR 2004)*. pp. 171–179.
19. Flouris, G., D. Plexousakis, and G. Antoniou: 2005, 'On Applying the AGM Theory to DLs and OWL.'. In: *Proceedings of the International Semantic Web Conference*. pp. 216–231.
20. Fuhrmann, A.: 1991, 'Theory Contraction Through Base Contraction'. *Journal of Philosophical Logic* **20**, 175–203.
21. Gärdenfors, P.: 1988, *Knowledge in Flux - Modeling the Dynamics of Epistemic States*. MIT Press.
22. Gärdenfors, P. and H. Rott: 1995, 'Belief Revision'. In: *Handbook of Logic in Artificial Intelligence and Logic Programming*, Vol. IV. Oxford University Press, Chapt. 4.2.
23. Grove, A.: 1988, 'Two Modellings for Theory Change'. *Journal of Philosophical Logic* **17**, 157–170.
24. Halaschek-Wiener, C. and Y. Katz: 2006, 'Belief Base Revision for Expressive Description Logics'. In: *Proceedings of the OWLED*06 Workshop on OWL: Experiences and Directions*.
25. Hansson, S. O.: 1992a, 'A dyadic representation of belief'. In: P. Gärdenfors (ed.): *Belief Revision*, Vol. 29 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, pp. 89–121.
26. Hansson, S. O.: 1992b, 'Reversing the Levi Identity'. *Journal of Philosophical Logic* **22**, 637–639.
27. Hansson, S. O.: 1994, 'Kernel Contraction'. *Journal of Symbolic Logic* **59**, 845–859.
28. Hansson, S. O.: 1999, *A Textbook of Belief Dynamics*. Kluwer Academic Press.
29. Hansson, S. O. and R. Wassermann: 2002, 'Local Change'. *Studia Logica* **70**(1), 49–76.
30. Jaśkowski, S.: 1969, 'Propositional Calculus for Contradictory Deductive Systems'. *Studia Logica* **24**, 143–157.
31. Katsuno, H. and A. O. Mendelzon: 1992, 'On the difference between updating a knowledge base and revising it'. In: P. Gärdenfors (ed.): *Belief Revision*, Vol. 29 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, pp. 183–203.
32. Lakemeyer, G. and W. Lang: 1996, 'Belief revision in a nonclassical logic'. In: G. Görz and S. Hölldobler (eds.): *KI-96: Advances in Artificial Intelligence*, Vol. 1137 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, pp. 199–211.
33. Levesque, H.: 1984, 'A logic of implicit and explicit belief'. In: *Proceedings of AAAI-84*.
34. Makinson, D.: 2009, 'Propositional relevance through letter-sharing'. *Journal of Applied Logic* **7**(4), 377 – 387. Special Issue: Formal Models of Belief Change in Rational Agents.
35. Mares, E. D.: 2002, 'A Paraconsistent Theory of Belief Revision'. *Erkenntnis* **56**(2), 229–246.

36. Nebel, B.: 1990, *Reasoning and Revision in Hybrid Representation Systems*, Vol. 422 of *Lecture Notes in Artificial Intelligence*. Springer-Verlag.
37. Parikh, R.: 1996, 'Beliefs, belief revision and splitting languages'. In: *Proceedings of Itallc-96*.
38. Parsia, B.: 2009, 'Topic-Sensitive Belief Revision'. Ph.D. thesis, University of Maryland.
39. Priest, G.: 2001, 'Paraconsistent belief revision'. *Theoria* **67**(3), 214 – 228.
40. Restall, G. and J. Slaney: 1995, 'Realistic belief revision'. In: M. D. Glas and Z. Pawlak (eds.): *Proceedings of the Second World Conference on the Fundamentals of Artificial Intelligence*. pp. 367–378.
41. Ribeiro, M. M. and R. Wassermann: 2006, 'First Steps Towards Revising Ontologies'. In: *Proceedings of the Second Workshop on Ontologies and their Applications (WONTO)*.
42. Ribeiro, M. M. and R. Wassermann: 2009a, 'AGM Revision in Description Logics'. In: *Proceedings the IJCAI Workshop on Automated Reasoning about Context and Ontology Evolution (ARCOE)*.
43. Ribeiro, M. M. and R. Wassermann: 2009b, 'Base Revision for Ontology Debugging'. *Journal of Logic and Computation* **19**(5), 721–743.
44. Ribeiro, M. M. and R. Wassermann: 2010, 'More About AGM Revision in Description Logics'. In: *Proceedings the ECAI Workshop on Automated Reasoning about Context and Ontology Evolution (ARCOE)*.
45. Rodrigues, O. T.: 1997, 'A Methodology for Iterated Information Change'. Ph.D. thesis, Imperial College, University of London.
46. Rott, H.: 2001, *Change, Choice and Inference: A Study of Belief Revision and Nonmonotonic Reasoning*. Oxford University Press.
47. Schaerf, M. and M. Cadoli: 1995, 'Tractable Reasoning via Approximation'. *Artificial Intelligence* **74**(2), 249–310.
48. Tanaka, K.: 2005, 'The AGM theory and inconsistent belief change'. *Logique et analyse* **48**(189-192), 113–150.
49. Wassermann, R.: 2001, 'On Structured Belief Bases'. In: H. Rott and M.-A. Williams (eds.): *Frontiers in Belief Revision*. Kluwer.
50. Wassermann, R.: 2003, 'Generalized Change and the Meaning of Rationality Postulates'. *Studia Logica* **73**(2), 299–319.