

Knowledge-base revision*

ODILE PAPINI

*SIS, équipe informatique, Faculté des Sciences et Techniques, Université de Toulon et du Var, BP 132, Avenue de l'université.
83957 La Garde Cedex, France. papini@univ-tln.fr*

Abstract

This paper surveys the most representative approaches of knowledge-base revision. After a description of the revision characterization according to the AGM paradigm, the paper reviews different revision methods. In each case, the same example is used, as a reference example, to illustrate the different approaches. Closely connected with revision, some other non-monotonic approaches, like update, are briefly presented.

1 Purpose

Modelling intelligent agent's reasoning requires designing knowledge bases for the purpose of performing symbolic reasoning. Among the different types of knowledge representation in the domain of artificial intelligence, logical representations stem from classical logic. However, this is not suitable for representing or treating items of information containing vagueness, incompleteness or uncertainty, or knowledge-base evolution that leads the agent to revise his beliefs about the world.

When a new item of information is added to a knowledge base, inconsistency can result. Revision means modifying the knowledge base in order to maintain consistency, while keeping the new information and removing the least possible previous information.

For example, referring to the famous Lea Sombe (1994), an agent believes that *Lea is a young woman*, that *she is a student*, and that *if Lea is a young woman and she is a student then she has no children*. But if he learns that *Lea has children*, inconsistency results, and consequently he has to revise his beliefs about Lea.

Because of the lack of information, the agent is induced to formulate assumptions, some of which could be invalidated when adding a new item of information. The conclusions drawn before the addition may not be valid anymore, and in fact, revisable reasoning is a special case of non-monotonic reasoning. In the example, if the agent assumes that Lea is a student and therefore has no children, if he learns that Lea has children, his assumption that she is a student could be not valid anymore.

The agent's beliefs and knowledge evolve with time and most applications in artificial intelligence, such as machine learning, planning, diagnosis, systems control or supervision, need a dynamic knowledge base for their representation, and require formalisation of revision operations. Revision plays a twofold role in planning: on the one hand the execution of actions involves revision of available knowledge about the world by the effects of actions; on the other hand the search for a plan involves a succession of revisions by sub-goals in order to reach the final goal. In diagnosis, the

*A preliminary draft of this paper was presented as invited conference at the 1996 *Journées Françaises de Programmation Logique avec Contraintes* (Papini 1996b). The final version of this paper was completed while the author was visiting the Simon Fraser University. Many thanks to J Doyle for helpful comments and to J Cohen, J Delgrande, D Dubois, JL Imbert, C Schwind and V Risch for helping the author to publish this survey.

revision of the intended description of a device by an observation or a sequence of observations can lead to a conflict which enables the detection of a failing component of the device.

As illustrated in the very simple example above, one of the problems arising when dealing with revision is that there are generally several possible ways to revise a knowledge base, and this has to be performed automatically. Because information is precious and must be preserved as much as possible, the principle of minimal change can provide a strategy to perform revision. On the other hand, practical implementations have to handle contradictory, uncertain or imprecise information. Several problems arise: how to represent knowledge in order to perform revision; how to define efficient revision operations (as revision generally can give several results); what result has to be chosen; and finally, according to a practical point of view, what computational model to support revision has to be provided?

Since knowledge-base revision is one of the main problems arising in knowledge representation, it has been tackled according to several points of view: symbolic or numerical, logical or probabilistic. This is both an old and important problem in artificial intelligence and a lot of work has been developed since 1980.

Although not exhaustive, the purpose of this paper is to review different aspects of revision, surveying the most important approaches with the same reference example. After an informal historical introduction in section 2, section 3 presents a formal framework for revision characterization; section 4 describes some syntactic approaches of knowledge base revision; sections 5 and 6 deal with semantic and mixed approaches respectively. The difference between revision and update is briefly presented in section 7; section 8 overviews numerical approaches of revision; revision and non-monotony are covered in section 9, the links between revision and constraints are examined in section 10; and finally section 11 gives an overview of actual research issues in the field, before concluding.

2 Introduction

The first works on revision were developed in the field of philosophical logic where they are often reported as theory change. The origin of the change theory logic comes from the works on consequence operations that Tarski put forward in the 1930s. A consequence operation, in the sense of Tarski, can be considered as a special case of the simplest kind of change, that is, the expansion operation. The expansion of a theory T^1 by a formula A is the set of the logical consequences of the union $T \cup \{A\}$.

In order to illustrate this, come back to the previous example slightly modified, and consider the following set of formulae expressed by means of propositions: *{Lea is a young woman; if she is a young woman then she is single; if she is a student then she has no children}*. The corresponding theory is the previous set of propositions with the additional consequence *Lea is single*. The expanded theory T by the formula $A = \{Lea is a student\}$ is the set of consequences of $T \cup \{A\} = \{Lea is a young woman; if she is a young woman then she is single; if she is a student then she has no children; Lea is single; Lea is a student; Lea has no children\}$.

Expansion is suitable when the added formula A is consistent with the initial theory T . However, if the formula $\neg A$ belongs to the theory T , the expansion of T by the formula A leads to an inconsistent theory. In order to define theory changes that preserve consistency, other operations have been introduced. The contraction operation allows the removal of a formula from a theory, whereas the revision operation allows the modification of the initial theory T in order to maintain consistency, keeping the added formula A .

The contraction of the previous theory T by the formula $A = \{Lea is single\}$ may give several results. For example, either *{Lea is a young woman; if she is a student then she has no children}* or *{if she is a young woman then she is single; if she is a student then she has no children}* or *{if Lea is a student then she has no children}* and so on. On the other hand, the revision of the theory T by the

¹A theory is a deductively closed set of formulae, expressed in given logical language.

formula $A = \{Lea \text{ is not single}\}$ may also give several results, for example, either $\{Lea \text{ is a young woman; if she is a student then she has no children; Lea is not single}\}$ or $\{if she is a young woman then she is single; if she is a student then she has no children; Lea is not single}\}$ or $\{if Lea is a student then she has no children; Lea is not single}\}$ and so on.

The first studies on change operations date from 1975 to 1977. They come from the works of Ellis and Levi (Levi 1980) in epistemology, and from the works of Harper (1975) in the history of science. Levi, in particular, defines the revision of a theory T by a formula A as the expansion by the formula A of the theory T first contracted by $\neg A$. This is called the Levi identity. In the previous example, revising the theory T by the formula $A = \{Lea \text{ is not single}\}$ will be performed by the contraction of the theory T by the formula $\neg A = \{Lea \text{ is single}\}$ and then by the expansion by the formula $A = \{Lea \text{ is not single}\}$.

The first formalisations of revision arose from philosophical logic with the works of Gärdenfors (1978), where revision is interpreted as belief change. Alchourron, Gärdenfors and Makinson (1985) formulated postulates, called the AGM (Alchourron *et al.* 1985) postulates, in order to characterize revision. These postulates stem from three principal ideas: (1) the consistency principle (a revision operation has to produce a consistent set of beliefs); (2) the principle of minimal change (a revision operation has to change the fewest possible beliefs); (3) a priority is given to the new item of information. These postulates focus on the logical structure of beliefs. They are based on the theory of consistency and do not take into account the justification of beliefs such as in TMS systems (Doyle 1979) mentioned below.

As the AGM postulates characterize but do not make it possible to construct revision operations, Alchourron, Gärdenfors and Makinson (1988) proposed a revision operation for a deductively closed set of formulae, called *partial meet contraction*. This operation uses the Levi identity which defines the revision operation from the contraction operation, i.e. the defined contraction operation involves the intersection of some maximal consistent subsets of formulae. They also defined a total order between formulae called *epistemic entrenchment*, which permits the principle of minimal change to hold.

On the other hand, truth maintenance systems have been developed in artificial intelligence, notably the TMS systems developed in 1979 by Doyle and the ATMS systems developed in 1986 by De Kleer which maintain the consistency of the knowledge base. In these systems, each elementary proposition is matched to its justification which represents the reasons to believe in the proposition, and the revision consists in giving up the propositions which have no more justification.

The question of how to perform update also arose in the field of databases. Fagin, Ullman and Vardi (Fagin *et al.*, 1983; Fagin *et al.*, 1986), particularly, proposed in 1983 a methodology for deductive database update based on model theory. Instead of a set of formulae, a set of interpretations of formulae is used to represent a knowledge base. In the above example, instead of the set of propositions $\{Lea \text{ is a young woman; if she is a young woman then she is single; if she is a student then she has no children}\}$, the set of interpretations² $\{\{Lea \text{ is a young woman, Lea is single, Lea is not a student, Lea has no children}\}$ and $\{Lea \text{ is a young woman, Lea is single, Lea is not a student, Lea has children}\}$ and $\{Lea \text{ is a young woman, Lea is single, Lea is a student, Lea has no children}\}$ is used for the sake of representing the knowledge base.

Significant links quickly appeared between works in philosophical logic, work developed in databases, and research performed in artificial intelligence. The collaboration of logicians and computer scientists has been fruitful and many works and results have ensued.

The various works performed on the subject can be classified according to different points of view: one can distinguish on the one hand syntactic and semantic approaches, and one can distinguish on the other hand update and revision.

In syntactic approaches, higher relevance is assigned to explicitly represented formulae, such as, for example, the revision operation of Nebel (1991), defined in the context of finite knowledge bases,

²An interpretation is denoted by a set of literals assigned true.

or Grove's work with systems of spheres (Grove 1988), both related to partial meet contraction drawn from Gärdenfors's works.

Semantic approaches stem from model theory, where a knowledge base is represented by only one formula. The revised knowledge-base models are the models of the added formula which are the closest to some models of the initial knowledge base. The principle of minimal change is defined in terms of distances or orders between models. This approach has been followed by several authors, for instance Borgida (1985), Dalal (1988) and Winslett (1988): these works only differ by the choice of the used order. These approaches were formalised in a common framework by Katsuno and Mendelzon (1991), who gave a new formulation of the AGM postulates, called KM formulation, and provided a representation theorem which shows an equivalence between these postulates and a revision process based on total pre-orders between interpretations. As an example, Spohn's *conditional ordinal function* illustrates a revision mechanism based on the ranking of interpretations (Spohn, 1987). Mixed approaches have also been proposed, in order to take advantage of both syntactic and semantic approaches (Papini & Rauzy, 1995; Willard & Yi, 1990).

Another consequence of the results of Katsuno and Mendelzon (1991a) was to get a better understanding regarding the distinction between revision and update. So let us consider a very simple example within the framework of actions, in which the knowledge base specifies that in a room both the window and the door may not be open. Suppose now that this knowledge base is revised by the action of somebody opening the door. A revision operation leads to the result that when somebody opens the door, the window is automatically closed, which is not desirable within the context of actions. Update operation, however, distinguishes two cases according to the fact that the window may or may not be closed before the execution of the action. Revising a knowledge base makes the knowledge-base models evolve as a whole towards the closest model of the new item of information. In contrast, updating a knowledge base makes each knowledge-base model evolve towards the closest model of the new item of information.

Other approaches have used numerical formalisms within the frameworks of both probability and possibility theories, such as, for example, Dubois, Lang and Prade (1994). These authors use degrees which express how the models may correspond to a real state of the world. The change of a state of knowledge by the introduction of a new item of information, resulting from a conditioning operation, leads to the modification of the distribution of probability, or possibility, respectively.

Since revision can also be regarded as a non-monotonic operation, another point of view was to define a non-monotonic formalism in order to represent revision reasoning, such as Reiter's default logic (Reiter 1980) and McCarthy's circumscription (McCarthy 1980a; 1980b). These formalisms are not described in the present paper, but the connection between non-monotonic formalisms and revision are presented in section 9.

Most approaches mentioned above deal with a one-step revision mechanism and are not suitable for an iteration of the revision process. As iterated revision is a central concern in AI applications, it has been the topic of recent works and is still a subject of ongoing research. This paper reviews the different aspects of revision mentioned above, from the first formulations of Gärdenfors to the most recent approaches dealing with iterated revision.

3 Revision characterisation

As mentioned in the introduction above, the first formalisations of the change operations come from the field of philosophical logic and involve the notion of theories. We now introduce more formally some notions useful for the following.

3.1 Preliminaries

An inference operation formalises the drawing of conclusions.

Definition 1 Let X be a set of formulae belonging to a classical logical language and let A be a formula,

$X \vdash A$ denotes that A can be inferred from X , and the inference relation is defined satisfying the following conditions:

- i) if $A \in X$ then $X \vdash A$ (reflexivity);
- ii) if $X \vdash A$ and $X \subseteq Y$ then $Y \vdash A$ (monotonicity);
- iii) if $X, A \vdash B$ and $Y \vdash A$ then $X, Y \vdash B$ (cut).

A set of logical conclusions can be defined as:

Definition 2 Let X be a set of formulae, $\text{Cons}(X)$ denotes the set of consequences of X , $A \in \text{Cons}(X)$ iff $X \vdash A$.

Consequence 3 If the inference relation $\vdash$ satisfies i), ii) and iii) then $\text{Cons}(X)$ satisfies:

- 1) $X \subseteq \text{Cons}(X)$;
- 2) $\text{Cons}(\text{Cons}(X)) \subseteq \text{Cons}(X)$;
- 3) if $X \subseteq Y$ then $\text{Cons}(X) \subseteq \text{Cons}(Y)$.

Conversly, if Cons satisfies 1), 2) and 3) then the inference relation $\vdash$ satisfies i), ii) and iii).

A theory is also called a deductively closed set of formulae, more formally:

Definition 4 Let X be a set of formulae, X is a theory iff $\text{Cons}(X) = X$.

3.2 Theory-change operations

The different theory-change operations mentioned in the previous section are now described more formally.

3.2.1 Theory expansion

Theory *expansion* corresponds to the introduction of a new formula into the theory, without modification of the initial theory.

Definition 5 Theory expansion. Let T be a theory and A be a formula, $T + A$ denotes the expansion of the theory T by a formula A . Since the result of an expansion is also closed under the consequence relation, notice that $T + A = \text{Cons}(T \cup \{A\})$.

Remark 1 From now on, the first example involving *Lea* will be used more formally, *young* denotes the proposition *Lea* is a young woman, *single* denotes the proposition *Lea* is single, and *has_no_children* denotes the proposition *Lea* has no children. The same example will be used several times within the rest of the paper as a reference example according to the various approaches of revision.

Example 1 Let T be the theory defined by $T = \text{Cons}(\{\text{young}, \text{single}, \text{young} \wedge \text{single} \rightarrow \text{has_no_children}\})$, the expansion of the theory by the formula *Lea* is a student denoted by *student* immediately gives

$$T + \text{student} = \text{Cons}(T \cup \{\text{student}\}),$$

thus

$$T + \text{student} = \text{Cons}(\{\text{young}, \text{single}, \text{young} \wedge \text{single} \rightarrow \text{has_no_children}, \text{has_no_children}, \text{student}\}).$$

Theory expansion is quite suitable when the added formula is consistent with the initial theory, whereas the introduction of a formula inconsistent with the initial theory leads to an inconsistent theory which cannot be dealt with by a classical logic. This is the reason why contraction and revision operations are defined.

3.2.2 Theory contraction

Theory *contraction* corresponds to the withdrawal of a formula from the theory. This operation is

more complex than it seems, because it can be performed in several ways, as shown in the example below.

Definition 6 Theory Contraction. *Let T be a theory and A be a formula, $T - A$ denotes the contraction of the theory T by the formula A . Since the result of a contraction is also closed under the consequence relation, note that $T - A = \text{Cons}(T - A)$.*

Example 2 *The contraction of the theory $T = \text{Cons}(\{\text{young}, \text{single}, \text{young} \wedge \text{single} \rightarrow \text{has_no_children}\})$ by the formula has_no_children can be performed in several ways:*

$T - \text{has_no_children}; = \text{Cons}(\{\text{young}, \text{single}\})$
or
 $T - \text{has_no_children} = \text{Cons}(\{\text{young}, \text{young} \wedge \text{single} \rightarrow \text{has_no_children}\})$
or
 $T - \text{has_no_children} = \text{Cons}(\{\text{single}, \text{young} \wedge \text{single} \rightarrow \text{has_no_children}\})$
or
 $T - \text{has_no_children} = \text{Cons}(\{\text{young}\})$ or...
or
 $T - \text{has_no_children} = \emptyset$

3.2.3 Theory revision

Theory revision corresponds to the addition of a formula into the theory, with the constraint that the resulting theory has to be consistent. This can lead to the removal of some formulae from the initial theory.

Definition 7 Theory revision. *Let T be a theory and let A be a formula, $T \star A$ denotes the revision of the theory T by a formula A . Since the result of a revision is also closed under the consequence relation, notice that $T \star A = \text{Cons}(T \star A)$.*

Example 3 *The revision of the theory $T = \text{Cons}(\{\text{young}, \text{single}, \text{young} \wedge \text{single} \rightarrow \text{has_no_children}\})$ by the formula $\neg(\text{has_no_children})$ can also be performed in several ways:*

$T \star \neg(\text{has_no_children}) = \text{Cons}(\{\text{young}, \text{single}, \neg(\text{has_no_children})\})$
or
 $T \star \neg(\text{has_no_children}) = \text{Cons}(\{\text{young}, \text{young} \wedge \text{single} \rightarrow \text{has_no_children}, \neg(\text{has_no_children})\})$
or
 $T \star \neg(\text{has_no_children}) = \text{Cons}(\{\text{single}, \text{young} \wedge \text{single} \rightarrow \text{has_no_children}, \neg(\text{has_no_children})\})$
or... or
 $T \star \neg(\text{has_no_children}) = \text{Cons}(\{\text{young}, \neg(\text{has_no_children})\})$
or
 $T \star \neg(\text{has_no_children}) = \text{Cons}(\{\neg(\text{has_no_children})\})$.

Like contraction, revision can be performed in several ways.

Remark 2 *As illustrated in the previous examples, the complexity of revision and contraction operations comes from the fact that these operations can be performed in several ways, and we do not know which way to select from them. However, information is precious, and one generally tries to remove the least information possible, which is illustrated by the principle of minimal change.*

The link between revision and contraction operations is established as follows:

Definition 8 identities. *Let T be a theory and let A be a formula, $-$ denotes a contraction operation and $\star$ denotes a revision operation.*

The Levi identity is defined as $T \star A = \text{Cons}((T - \neg A) \cup \{A\})$.

The Harper identity is defined as $T - A = \text{Cons}(T \cap (T \star \neg A))$.

3.3 AGM postulates

Alchourron, Gärdenfors and Makinson (1985) proposed a formal framework in which revision is interpreted as belief change. Focusing on the logical structure of beliefs, they formulate eight postulates which a revised theory has to verify. These postulates stem from three main principles: the new item of information has to appear in the revised theory, the revised theory has to be consistent, and revision operation has to change the fewest possible beliefs.

AGM postulates Let T be a theory, and let A and B be formulae. $T \star A$ denotes the theory T revised by A . $T + A$ is the smallest deductively closed set of formulae containing both T and A . $T^\perp$ denotes the set of all the formulae.

- (G★1) $T \star A$ is a theory.
- (G★2) $A \in T \star A$.
- (G★3) $T \star A \subseteq T + A$.
- (G★4) If $\neg A \notin T$ then $T \star A = T + A$.
- (G★5) $T \star A = T^\perp$ only if A is unsatisfiable.
- (G★6) If $A \equiv B$ then $T \star A = T \star B$.
- (G★7) $T \star (A \wedge B) \subseteq (T \star A) + B$.
- (G★8) If $\neg B \notin T \star A$ then
 $(T \star A) + B = T \star (A \wedge B)$.

The postulate (G★1) expresses the fact that a theory revised by a formula is a theory. (G★2) specifies that the formula A belongs to the revised theory. (G★3) and (G★4) give the result of the revision when A is consistent with T . (G★5) is linked to the preservation and the restoration of consistency; (G★6) specifies that the result of the revision has to be independent from the syntactic form of the added formula. (G★7) and (G★8) point that when A is consistent with T the change has to be minimal. Minimally revising T to include both A and B should reduce to an expansion of $T \star A$, so long as B does not contradict $T \star A$.

The eight postulates that a contracted theory has to verify are established in the same way. $T - A$ denotes the theory T contracted by A .

- (G-1) $T - A$ is a theory.
- (G-2) $T - A \subseteq T$.
- (G-3) If $A \notin T$ then $T - A = T$.
- (G-4) If $\not\models A$ then $A \notin T - A$.
- (G-5) $T \subseteq ((T - A) + A)$.
- (G-6) If $A \equiv B$ then $T - A = T - B$.
- (G-7) $(T - A) \cap (T - B) \subseteq T - (A \wedge B)$.
- (G-8) If $A \notin T - (A \wedge B)$ then
 $T - (A \wedge B) \subseteq T - A$.

3.4 Revision operations proposed for theories

AGM postulates are non-constructive; they do not give a method for the construction of revision operations. After these formulations, Alchourron, Gärdenfors and Makinson (Gärdenfors, 1988) proposed revision operations for theories. Given a contraction operation, a related revision operation can be defined via the Levi identity. The construction of a contraction operation involves the determination of maximal consistent subsets not implying a formula $\neg A$. In the following, $T \downarrow \neg A$ denotes the set of these subsets such that:

$T \downarrow \neg A = \{T' \mid T' \subseteq T, T' \not\models \neg A\}$ and
 if $T' \subset M \subseteq T$ then $M \models \neg A$.

3.4.1 Full meet contraction

The first contraction operation, called *full meet contraction*, involves the intersection between all the maximal consistent subsets which fail to imply $\neg A$. More formally:

$$T - \neg A = \begin{cases} \cap T \downarrow \neg A, & \text{if } T \downarrow \neg A \text{ is not empty;} \\ T & \text{otherwise.} \end{cases}$$

This contraction operation satisfies the first six AGM postulates, i.e. (G-1) to (G-6). However, this operation is too cautious to define a reasonable revision operation: if $\neg A \in T$ and $T - \neg A$ is defined by this operation then $T - \neg A = \{B \in T, A \vdash B\}$. Consequently, the revision operation provided by the Levi identity is obvious, that is $T \star A = \text{Cons}(\{A\})$. For more details see Sombe (1994).

Example 4 Let T be the theory $T = \text{Cons}(\{\text{young}, \text{single}, \text{young} \wedge \text{single} \rightarrow \text{has_no_children}\})$. There are 4 maximal consistent subsets which fail to imply $\neg(\text{has_no_children})$:

$$T \downarrow \text{has_no_children} = \{\text{Cons}(\{\text{young}, \text{single}\}), \text{Cons}(\{\text{young}, \text{single} \leftrightarrow \text{has_no_children}\}), \text{Cons}(\{\text{single}, \text{young} \leftrightarrow \text{has_no_children}\}), \text{Cons}(\{\text{young} \leftrightarrow \text{has_no_children}, \text{single} \leftrightarrow \text{has_no_children}\})\}$$

The full meet contraction operation gives the intersection between the maximal consistent subsets:

$$T - \text{has_no_children} = \cap \{T \downarrow \text{has_no_children}\} = \emptyset.$$

The revision operation provided by the Levi identity is $T \star \neg(\text{has_no_children}) = \text{Cons}(\{\neg(\text{has_no_children})\})$. This means that we remove all previous information which is not reasonable according to the principle of minimal change.

3.4.2 Maxi choice contraction

Another way to construct a contraction operation is to choose one of the maximal consistent subsets which fail to imply $\neg A$, according to a selection function s . More formally:

$$T - \neg A = \begin{cases} s(T \downarrow \neg A), & \text{if } T \downarrow \neg A \text{ is not empty;} \\ T & \text{otherwise.} \end{cases}$$

This contraction operation, called *maxi choice contraction*, satisfies the first six AGM postulates, i.e. (G-1) to (G-6). However, this operation is too brave to define a reasonable revision operation: if $\neg A \in T$ and $T - \neg A$ is defined by this operation then for each B , either $\neg A \vee B \in T - \neg A$ or $\neg A \vee \neg B \in T - \neg A$. Consequently, by the Levi identity, for each B , either $B \in T \star A$ or $\neg B \in T \star A$. For more details see Sombe (1994).

Example 5 Let T be the theory $T = \text{Cons}(\{\text{young}, \text{single}, \text{young} \wedge \text{single} \rightarrow \text{has_no_children}\})$. There are 4 maximal consistent subsets which fail to imply $\neg(\text{has_no_children})$:

$$T \downarrow \text{has_no_children} = \{\text{Cons}(\{\text{young}, \text{single}\}), \text{Cons}(\{\text{young}, \text{single} \leftrightarrow \text{has_no_children}\}),$$

$$\text{Cons}(\{\text{single}, \text{young} \leftrightarrow \text{has_no_children}\}), \text{Cons}(\{\text{young} \leftrightarrow \text{has_no_children}, \text{single} \leftrightarrow \text{has_no_children}\})\}.$$

The maxi choice contraction operation gives one of the maximal consistent subsets according to the selection function s :

$$T - \text{has_no_children} = s(T \downarrow \text{has_no_children}).$$

which yields:

$$\text{Cons}(\{\text{young}, \text{single}\})$$

or

$$\text{Cons}(\{\text{young}, \text{single} \leftrightarrow \text{has_no_children}\})$$

or

$Cons(\{single, young \leftrightarrow has_no_children\})$

or

$Cons(\{young \leftrightarrow has_no_children, single \leftrightarrow has_no_children\})$

Suppose we have no idea about the fact that Lea is student or not, and that we believe T , if we revise T by $\neg(has_no_children)$ using a revision operation provided by Levi identity, we have to believe that Lea is either a student or not, which is not reasonable.

3.4.3 Partial meet contraction

Attempting to define a contraction operation that could be a compromise between the two previous operations seems natural. This is performed by the operation which involves the intersection of some maximal consistent subsets which fail to imply $\neg A$, chosen according to a selection function s . More formally :

$$T - \neg A = \begin{cases} \bigcap s(T \downarrow \neg A), & \text{if } T \downarrow \neg A \text{ is not empty;} \\ T & \text{otherwise.} \end{cases}$$

This contraction operation, called *partial meet contraction*, satisfies the first six AGM postulates, i.e. $(G-1)$ to $(G-6)$, it gives satisfying results. For more details see Sombe (1994).

Example 6 Let T be the theory $T = Cons(\{young, single, young \wedge single \rightarrow has_no_children\})$. There are 4 maximal consistent subsets which fail to imply $\neg(has_no_children)$:

$T \downarrow has_no_children = \{Cons(\{young, single\}),$
 $Cons(\{young, single \leftrightarrow has_no_children\}),$
 $Cons(\{single, young \leftrightarrow has_no_children\}),$
 $Cons(\{young \leftrightarrow has_no_children, single \leftrightarrow has_no_children\})\}.$

The partial meet contraction operation gives the intersection of some maximal consistent subsets:

$$T - has_no_children = \bigcap s(T \downarrow has_no_children).$$

There are 15 possibilities:

$Cons(\{young, single\})$

or

$Cons(\{young, single \leftrightarrow has_no_children\})$

or

$Cons(\{single, young \leftrightarrow has_no_children\})$

or

$Cons(\{young \leftrightarrow has_no_children, single \leftrightarrow has_no_children\})$

or

$Cons(\{young, has_no_children \rightarrow single\})$

or

$Cons(\{single, has_no_children \rightarrow young\})$

or

$Cons(\{has_no_children \rightarrow young, has_no_children \rightarrow single, has_no_children \vee (young \leftrightarrow single)\})$

or

$Cons(\{young \vee single, young \wedge single \rightarrow has_no_children, has_no_children \rightarrow young, has_no_children \rightarrow single\})$

or

$Cons(\{has_no_children \rightarrow young, single \leftrightarrow has_no_children\})$

or

$Cons(\{young \leftrightarrow has_no_children, has_no_children \rightarrow single\})$

or

$Cons(\{young \vee single, has_no_children \rightarrow young, has_no_children \rightarrow single\})$

or
 $Cons(\{single \rightarrow young \vee has_no_children, has_no_children \rightarrow young, has_no_children \rightarrow single\})$
 or
 $Cons(\{young \rightarrow single \vee has_no_children, has_no_children \rightarrow young, has_no_children \rightarrow single\})$
 or
 $Cons(\{young \wedge single \rightarrow has_no_children, has_no_children \rightarrow young, has_no_children \rightarrow single\})$
 or
 $Cons(\{has_no_children \rightarrow young, has_no_children \rightarrow single\})$.

The four first possibilities give the same result as the maxi choice contraction. So the choice of the contracted theory is between the eleven following possibilities; this choice has to be performed according the principle of minimal change and the revision operations provided by the Levi identity give reasonable results.

3.5 Epistemic entrenchment

As shown in the previous example, the contraction operation may yield several solutions, that is, several maximal consistent subsets. The choice of one set among the others has to be performed according to the principle of minimal change. The partial meet contraction operation enables us to carry out this choice, although it is rather difficult to achieve since the operation involves the comparison of several subsets of formulae.

Gärdenfors has proposed defining an ordering relation among formulae, following philosophical considerations. For example, in the scientific theories revision, when a contradiction occurs, not all the formulae are equally involved. Some may be more easily revised than others; they are less *entrenched*.

Definition 9 Epistemic entrenchment. *The epistemic entrenchment is an order relation, denoted $\leq_{EE}$, such that $\forall A, B \in T, A \leq_{EE} B$ iff A is less entrenched than B .*

This order relation has to verify the following property:

- 1) The relation $\leq_{EE}$ is transitive.
- 2) $\forall A, B \in T$ if $A \vdash B$ then $A \leq_{EE} B$.
- 3) $\forall A, B \in T, A \leq_{EE} A \wedge B$ or $B \leq_{EE} A \wedge B$.

Let $\leq_{EE}$ be an epistemic entrenchment relation, the revision $T \star A$ is carried out in two steps. First, a maximal consistent subset which fails to imply $\neg A$ is constructed following the epistemic entrenchment relation, i.e. $B \in T - \neg A$ iff $\neg A <_{EE} B$. In other words, only the formulae more entrenched than $\neg A$ have to be kept ($A <_{EE} B$ iff $A \leq_{EE} B$ and not $B \leq_{EE} A$), the Levi identity being then applied. Gärdenfors and Makinson showed that each operation verifying the AGM postulates can be expressed by means of an epistemic entrenchment relation.

These works are of a great interest, but nevertheless they remain abstract and are non-constructive regarding the determination of selection functions or preference relations. Moreover, these methods propose too many solutions to be practically implemented. Furthermore, there is no mean to specify how an epistemic entrenchment ordering can be obtained in case of iterated revision.

4 Syntactic approaches

When knowledge is represented by a deductively closed set of formulae, all the formulae are considered at the same level. However, it would be interesting to distinguish explicitly given formulae from derived ones. Syntactic approaches to revision give preference to explicitly given formulae and define a hierarchy.

4.1 Finite knowledge-base revision

Drawing inspiration from theory revision, Nebel (1991) first proposed a revision operation for finite knowledge bases as follows. Let B be a finite set of formulae, representing a knowledge base, and x be a formula. $B \downarrow x$ denotes the set of maximal consistent subsets which fail to imply x :

$B \downarrow x = \{B', B' \subseteq B, B' \not\models x \text{ and if } B' \subset M \subseteq B \text{ then } M \vdash x\}$, the contraction operation is defined by:

$$B - x = \begin{cases} (\bigvee B')_{B' \in B \downarrow x}, & \text{if } \not\models x; \\ B & \text{otherwise.} \end{cases}$$

The revision operation is then given by:

$$B \oplus x = (B - \neg x) \wedge x.$$

Gärdenfors's postulates deal with deductively closed sets of formulae; they do not directly apply to finite knowledge bases. However, let K be the deductive closure of B . Both the contraction operation defined on $K - x$ and the deductive closure of $B - x$ satisfy the postulates (G-1) to (G-4) and (G-6) but do not satisfy (G-5).

Example 7 Let B be a finite knowledge base $B = \{\text{young}, \text{young} \rightarrow \text{single}\}$, then $K = \text{Cons}(\{\text{young}, \text{single}\})$. If B is contracted by single , $B \downarrow \text{single} = \{\{\text{young}\}, \{\text{young} \rightarrow \text{single}\}\}$. Therefore, according to the definition of the contraction operation, $B - \text{single} = \text{young} - (\text{young} \rightarrow \text{single})$ which is a tautology and $((K - \text{single}) + \text{single})$ equals $\text{Cons}(\{\text{single}\})$.

As a result $K \not\subseteq ((K - \text{single}) + \text{single})$.

In order to satisfy the postulate (G-5), Nebel then redefined the contraction operation as follows:

$$B - x = \begin{cases} (\bigvee B')_{B' \in B \downarrow x} \wedge (B \vee \neg x), & \text{if } \not\models x; \\ B & \text{otherwise.} \end{cases}$$

In the revision operations defined above, all the knowledge-base formulae are considered at the same level. In order to satisfy the principle of minimal change, Nebel stratifies the knowledge base in disjoint priority classes, $B_i, i \geq 1$, such that the formulae of B_i have higher priority than those of B_j iff $i < j$. $B \Downarrow x$ denotes the set of maximal consistent subsets which fail to imply x according to the priorities. The prioritized revision operation is then defined as follows:

$$B \oplus^p x = \text{Cons}((\bigvee (B \Downarrow \neg x)) \wedge x).$$

The interest of this method lies in the fact that it proposes an effective revision operation which leads to satisfying results and also verifies the AGM postulates. However, the finite knowledge base is sensitive to syntactic change. Moreover, the revised base is expressed by a disjunction which renders difficult the iteration of the revision process.

Example 8 Let B be a finite knowledge base, $B = \{\text{young}, \text{single}, \text{young} \wedge \text{single} \rightarrow \text{has_no_children}\}$.

There are 3 maximal consistent subsets:

$B \downarrow \text{has_no_children} =$
 $\{\{\text{young}, \text{single}\},$
 $\{\text{young}, \text{young} \wedge \text{single} \rightarrow \text{has_no_children}\},$
 $\{\text{single}, \text{young} \wedge \text{single} \rightarrow \text{has_no_children}\}\}.$

The finite subsets in $B \downarrow \text{has_no_children}$ are taken as the conjunction of the formulae appearing in them since they are used as formulae in $B - \text{has_no_children}$. Therefore, the contracted finite base B by has_no_children is:

$$B - \text{has_no_children} = (\text{young} \wedge \text{single}) \vee (\text{young} \wedge (\text{young} \wedge \text{single} \rightarrow \text{has_no_children})) \vee (\text{single} \wedge (\text{young} \wedge \text{single} \rightarrow \text{has_no_children})).$$

It can easily be shown that

$$B - \text{has_no_children} \Leftrightarrow \text{young} \vee \text{single}.$$

On the other hand,

$$B \vee \neg(\text{has_no_children}) = (\text{young} \wedge \text{single} \wedge (\text{young} \wedge \text{single} \rightarrow \text{has_no_children})) \vee \neg(\text{has_no_children}).$$

Now,

$$(\text{young} \wedge \text{single} \wedge (\text{young} \wedge \text{single} \rightarrow \text{has_no_children})) \vee \neg(\text{has_no_children}) \Leftrightarrow \text{young} \wedge \text{single} \wedge \text{has_no_children},$$

therefore

$$B \vee \neg(\text{has_no_children}) = (\text{young} \wedge \text{single} \wedge \text{has_no_children}) \vee \neg(\text{has_no_children}) \text{ and} \\ B \vee \neg(\text{has_no_children}) \Leftrightarrow \text{has_no_children} \rightarrow \text{young} \wedge \text{single} \wedge \text{has_no_children}$$

thus

$$B \oplus \neg(\text{has_no_children}) = (B - \text{has_no_children}) \wedge (B \vee \neg(\text{has_no_children})) \wedge \neg(\text{has_no_children}) \\ = (\text{young} \vee \text{single}) \wedge (\text{has_no_children} \rightarrow \text{young} \wedge \text{single} \wedge \text{has_no_children}) \wedge \neg(\text{has_no_children}) \\ = (\text{young} \vee \text{single}) \wedge \neg(\text{has_no_children}).$$

Revising by $\neg(\text{has_no_children})$ we cannot any more have both *young* and *single* in the revised finite knowledge base. Nebel's revision operation keeps $\text{young} \vee \text{single}$, which means that we have $\text{young} \wedge \neg(\text{has_no_children})$ or $\text{single} \wedge \neg(\text{has_no_children})$ in the revised finite knowledge base, which makes sense and follows the principle of minimal change.

4.2 Revision according to the syntactic possible worlds approach

Ginsberg and Smith (1988), within the framework of actions, define knowledge bases as sets of formulae called worlds. The introduction of a new formula into a world yields the closest world to that of the added formula. The notion of closeness between worlds is defined by set inclusion: let W , W_1 and W_2 be worlds, W_1 is closer to W than W_2 if the formulae of W_2 that belong to W also belong to W_1 .

Different kinds of formulae can be distinguished: P , the set of *protected* formulae, true in every world, and Q , the set of *qualifications*, representing constraints which, if they are not satisfied, prohibit the action to be performed.

In case of lack of qualification, the introduction of a set of formulae C into a world S leads to the definition of a *potential* world for C in S , i.e. every subset of $S \cup C$ which is consistent and contains both C and P . A *possible* world for C in S is a potential world which is maximal for set inclusion.

In case of presence of qualifications, the introduction of a set of formulae C into a world S leads to the construction of possible worlds for C in S . The possible worlds are computed as if the set of protected formulae was $P - Q$, and possible worlds inconsistent with Q are removed.

The introduction of a set of formulae C into a world S entails that several possible worlds may result. In this case Ginsberg and Smith suggest that either knowledge is incomplete, for instance a physical law is missing, or the worlds have to be ordered according to an order defined on the basic facts of the world. However, none of these solutions seem to be fully satisfactory.

Example 9 Consider the previous example.

Without qualification:

Let S be a world, $S = \{young, single, young \wedge single \rightarrow has_no_children\}$ and C be the added set of formulae, $C = \{\neg(has_no_children)\}$. There are 3 possible worlds:

$W_1 = \{young, single, \neg(has_no_children)\}$,

$W_2 = \{young, young \wedge single \rightarrow has_no_children, \neg(has_no_children)\}$ and

$W_3 = \{single, young \wedge single \rightarrow has_no_children, \neg(has_no_children)\}$.

However, in order to select between these worlds, we need some extra information, like an ordering between worlds.

In presence of qualifications:

Let S be a world, $S = \{young, single, \}$ and Q be a qualification, $Q = \{young \wedge single \rightarrow has_no_children\}$ and C be the added set of formulae, $C = \{\neg(has_no_children)\}$. The only possible world is $W = \{young, single, \neg(has_no_children)\}$, which is inconsistent with Q . Since the formulae $young$ and $single$ belong to S , they are preferred in the possible world after revision. This revision operation removes $\neg(has_no_children)$ which belongs to Q and follows the principle of minimal change, since it only removes one formula.

4.3 Revision according to Grove's system of spheres

Grove (1988) proposed an alternative approach of revision using a system of spheres based on the semantics of spheres defined by Lewis (1973). In order to define a revision operation, let $\mathcal{L}$ be a logical language. Grove focuses on the set of maximal consistent subsets of $\mathcal{L}$, denoted $\mathcal{M}$. The elements of $\mathcal{M}$ are also called possible worlds. Within this framework, a knowledge base can be represented by a set of formulae of $\mathcal{L}$, denoted $[K]$, which is a subset of $\mathcal{M}$ and which consists of all the maximal consistent subsets containing the formulae of K . More formally,

$$[K] = \{M \in \mathcal{M} \mid K \subseteq M\}.$$

A system of spheres is now defined as follows:

Definition 10 A system of spheres centered on $[K]$ is a collection of subsets of $\mathcal{M}$, denoted $\mathcal{S}$, which satisfies the following conditions:

- 1) $\mathcal{S}$ is totally ordered by $\subseteq$.
- 2) $[K]$ is a $\subseteq$ -minimum of $\mathcal{S}$.
- 3) $\mathcal{M}$ is in $\mathcal{S}$.
- 4) If A is a formula of $\mathcal{L}$ and there is a sphere intersecting $[A]$, then there is a smallest sphere in $\mathcal{S}$ intersecting $[A]$.

In other words, condition 1) says that if S and S' are in $\mathcal{S}$, then $S \subseteq S'$ or $S' \subseteq S$. Condition 2) means that $[K] \in \mathcal{S}$ and if $S \in \mathcal{S}$, then $[K] \subseteq S$.

Within the framework of the system of spheres, revision can be performed as follows. Let S_A be the smallest sphere of $\mathcal{S}$ which intersects with $[A]$. The revision of the knowledge base K by a formula A can be represented by $[A] \cap S_A$. In fact, $[A] \cap S_A$ is the set of *closest* elements in $\mathcal{M}$ to $[K]$ which contain A . This representation of revision is quite appropriate: the defined revision operation, denoted $K \star A = K_{c(A)}$ satisfies the eight AGM postulates. Conversely for any revision operation $\star$ satisfying the eight AGM postulates, then for any knowledge base K there exists a system of spheres $\mathcal{S}$ centered on $[K]$ that satisfies $K \star A = K_{c(A)}$.

Grove also introduces a revision operation based on the ordering of the formulae of $\mathcal{L}$, which determines which formulae should be retained in $K \star A$. This ordering, denoted $\leq$, has to satisfy the following conditions:

Let A, B, C be formulae of $\mathcal{L}$ and let K be a knowledge base,

- $\leq_1$) $\leq$ is connected: $A \leq B$ or $B \leq A$, for all A, B in $\mathcal{L}$.
- $\leq_2$) $\leq$ is transitive: $A \leq B$ and $B \leq C$ implies $A \leq C$.
- $\leq_3$) If $A \rightarrow B - C$ is in $\mathcal{L}$ then either $B \leq A$ or $C \leq A$.
- $\leq_4$) $\neg A \notin K$ iff $A \leq B$, for all B in $\mathcal{L}$.
- $\leq_5$) $\neg A \in L$ iff $B \leq A$, for all B in $\mathcal{L}$.

This ordering may be used to define a revision operation:

$$(\bullet) \quad B \in K \star A \text{ iff } A \wedge B < A \wedge \neg B.$$

Grove (1988) has shown that for any revision operation satisfying the eight AGM postulates, for any knowledge base K there exists a relation $\leq$ such that $(\bullet)$ and $\leq_1 - \leq_5$ are satisfied. Conversely, let a relation satisfy $\leq_1 - \leq_5$, then there is a corresponding revision operation defined as $(\bullet)$ which satisfies the eight AGM postulates.

As the system of spheres has been used as a logical framework for conditionals, this approach is of great interest, because it was the first to make the connection between revision and conditionals which are to be taken into account when dealing with iterated revision. For more details see section 11.

Example 10 Let K be a finite knowledge base represented by the set of formulae,

$$K = \{\text{young, single, young} \wedge \text{single} \rightarrow \text{has_no_children}\}.$$

In this approach, the knowledge base is represented by: $[K] = \{\text{young, single, young} \wedge \text{single} \rightarrow \text{has_no_children}\}$.

Let $\neg(\text{has_no_children})$ be the added formula; the set consisting of all maximal consistent subsets containing $\neg(\text{has_no_children})$ is:

$$\begin{aligned} [\neg(\text{has_no_children})] = \\ \{ \{ \neg(\text{has_no_children}), \text{young, single} \}, \\ \{ \neg(\text{has_no_children}), \text{young, young} \wedge \text{single} \rightarrow \text{has_no_children} \}, \\ \{ \neg(\text{has_no_children}), \text{single, young} \wedge \text{single} \rightarrow \text{has_no_children} \} \}. \end{aligned}$$

The smallest sphere intersecting with $\neg(\text{has_no_children})$, denoted $S_{\neg(\text{has_no_children})}$ is:

$$\begin{aligned} S_{\neg(\text{has_no_children})} = \\ \{ \{ \text{young, single, young} \wedge \text{single} \rightarrow \text{has_no_children} \}, \\ \{ \neg(\text{has_no_children}), \text{young, single} \}, \\ \{ \neg(\text{has_no_children}), \text{young, young} \wedge \text{single} \rightarrow \text{has_no_children} \}, \\ \{ \neg(\text{has_no_children}), \text{single, young} \wedge \text{single} \rightarrow \text{has_no_children} \} \} \end{aligned}$$

and the revision of K by $\neg(\text{has_no_children})$ is:

$$\begin{aligned} [\neg(\text{has_no_children})] \cap S_{\neg(\text{has_no_children})} = \\ \{ \{ \neg(\text{has_no_children}), \text{young, single} \}, \\ \{ \neg(\text{has_no_children}), \text{young, young} \wedge \text{single} \rightarrow \text{has_no_children} \}, \\ \{ \neg(\text{has_no_children}), \text{single, young} \wedge \text{single} \rightarrow \text{has_no_children} \} \}. \end{aligned}$$

This set contains three subsets; we have to choose one of them. There is no information about how to make this choice. However, the result of revision follows the principle of minimal change since one formula is removed from each subset of the initial knowledge base.

5 Semantic approaches

Within the framework of semantic approaches, the knowledge base is represented by only one propositional formula ψ . The revision operation, denoted $\psi \circ \mu$, consists in finding models of μ

which are the closest possible (in the chosen metric) to models of ψ . These approaches stem from the principle of irrelevance of syntax. This principle asserts that the knowledge base resulting from a revision operation must to be independent of the syntax of the original knowledge base as well as independent of the syntax of the revision itself.³

5.1 Reformulation of AGM postulates

Katsuno and Mendelzon (1991b) unified the semantic approaches in a common framework in which they reformulated the AGM postulates. If a knowledge base is represented on one hand by K , a set of formulae, and on the other hand by ψ , a propositional formula such that $K = \{\phi \mid \psi \vdash \phi\}$, then a correspondence between $K \star \mu$ and $\phi \circ \mu$ has been established as follows:

KM formulation of AGM postulates Let ψ , ϕ and μ be formulae,

- (R1) $\psi \circ \mu$ implies μ .
- (R2) If $\psi \wedge \mu$ is satisfiable, then $\psi \circ \mu \equiv \psi \wedge \mu$.
- (R3) If μ is satisfiable, then so is $\psi \circ \mu$.
- (R4) If $\psi_1 \equiv \psi_2$ and $\mu_1 \equiv \mu_2$, then
 $\psi_1 \circ \mu_1 \equiv \psi_2 \circ \mu_2$.
- (R5) $(\psi \circ \mu) \wedge \phi$ implies $\psi \circ (\mu \wedge \phi)$.
- (R6) If $(\psi \circ \mu) \wedge \phi$ is satisfiable, then
 $\psi \circ (\mu \wedge \phi)$ implies $(\psi \circ \mu) \wedge \phi$.

(R1) specifies that the added formula belongs to the revised knowledge base, (R2) gives the revised knowledge base when the added formula is consistent with the initial knowledge base, (R3) ensures that no inconsistency is introduced in the revised knowledge base, (R4) expresses the principle of irrelevance of syntax, (R5) and (R6) are the direct translation of both $(G \star 7)$ and $(G \star 8)$ postulates.

5.2 Principle of minimal change

The principle of minimal change leads to the definition of orders between interpretations. Katsuno and Mendelzon synthesised the different proposed metrics allowing both comparison of the different interpretations and keeping most of the above postulates.

Let I be the set of all the interpretations and $Mod(\psi)$ be the set of models of ψ . A pre-order on I , denoted $\leq_\psi$, is linked with ψ . The relation $<_\psi$ is defined from $\leq_\psi$ as usual:

$$I <_\psi I' \text{ iff } I \leq_\psi I' \text{ and } I' \not\leq_\psi I.$$

The pre-order $\leq_\psi$ is *faithful* to ψ if it verifies the following conditions:

- 1) if $I, I' \in Mod(\psi)$ then $I <_\psi I'$ does not hold;
- 2) if $I \in Mod(\psi)$ and $I' \notin Mod(\psi)$ then $I <_\psi I'$ holds;
- 3) if $\psi \equiv \phi$ then $\leq_\psi = \leq_\phi$.

A minimal interpretation may thus be defined by:

$\mathcal{M} \subseteq I$, the set of minimal interpretations in $\mathcal{M}$ according to $\leq_\psi$ is denoted $\text{Min}(\mathcal{M}, \leq_\psi)$. And I is minimal in $\mathcal{M}$ according to $\leq_\psi$, if $I \in \mathcal{M}$ and there is no $I' \in \mathcal{M}$ such that $I' <_\psi I$.

Let ψ be a formula representing a knowledge base. Katsuno and Mendelzon showed that a revision operation satisfies the postulates (R1)–(R6) if and only if there exists a total pre-order $\leq_\psi$ such that $Mod(\psi \circ \mu) = \text{Min}(Mod(\mu), \leq_\psi)$.

³Not every author agrees with this principle, which is illustrated by the famous restaurant example proposed by Hansson. See Katsuno & Mendelzon (1991b).

5.2.1 Revision according to Dalal

The revision operation defined by Dalal (1988) uses the Hamming distance as metric. The Hamming distance between two interpretations is the number of propositional variables on which two interpretations differ. More formally, let I, J be two interpretations, the Hamming distance between I and J is denoted $d(I, J)$. The distance between $Mod(\psi)$ and I is defined by:

$$d(Mod(\psi), I) = \min_{J \in Mod(\psi)} d(J, I)^4$$

and given the notations of the previous:

$$I \leq_{\psi} J \text{ iff } d(Mod(\psi), I) \leq d(Mod(\psi), J).$$

The revision operation defined by Dalal may be written:

$$Mod(\psi \circ_D \mu) = \min(Mod(\mu), \leq_{\psi}).$$

This revision operation satisfies the postulates (R1)–(R6).

Example 11 Let ψ be the unique formula representing the knowledge base $\psi = \text{young} \wedge \text{single} \wedge (\text{young} \wedge \text{single} \rightarrow \text{has_no_children})$ and let μ be the added formula $\mu = \neg(\text{has_no_children})$. The only model of ψ is

$$I = \{\text{young}, \text{single}, \text{has_no_children}\}.$$

The models of μ are:

$$\begin{aligned} J_1 &= \{\text{young}, \text{single}, \neg(\text{has_no_children})\}, \\ J_2 &= \{\text{young}, \neg(\text{single}), \neg(\text{has_no_children})\}, \\ J_3 &= \{\neg(\text{young}), \text{single}, \neg(\text{has_no_children})\} \text{ and} \\ J_4 &= \{\neg(\text{young}), \neg(\text{single}), \neg(\text{has_no_children})\}. \end{aligned}$$

Using the Hamming distance,

$$d(I, J_1) = 1, d(I, J_2) = 2, d(I, J_3) = 2, d(I, J_4) = 3,$$

consequently,

$$J_1 = \{\text{young}, \text{single}, \neg(\text{has_no_children})\} \text{ is the model of } \psi \circ_D \mu.$$

In Dalal's revision operation, the principle of minimal change takes the form of the minimal number of the propositional variables which change. Since the added formula is $\neg(\text{has_no_children})$, and the model of ψ is $I = \{\text{young}, \text{single}, \text{has_no_children}\}$, the model of μ which only differs by one propositional variable is $J_1 = \{\text{young}, \text{single}, \neg(\text{has_no_children})\}$.

5.2.2 Revision according to A. Borgida

The revision operation defined by Borgida (1985) focuses on sets of propositional variables on which two models differ. Let I, J be two interpretations, $\text{diff}(I, J)$ denotes the set of propositional variables on which the interpretations I and J differ. Let $\text{diff}(I, \mu) = \{\text{diff}(I, J)\}$, where J is a model of μ , the revision operation according to Borgida is defined as follows:

if μ is consistent with ψ then $\psi \circ_B \mu = \psi \wedge \mu$;
else $Mod(\psi \circ_B \mu) = \min_{I \in Mod(\psi)} (\text{diff}(I, \mu))$.

This revision operation satisfies the postulates (R1)–(R5).

Example 12 Let ψ be the unique formula representing the knowledge base $\psi = \text{young} \wedge \text{single} \wedge (\text{young} \wedge \text{single} \rightarrow \text{has_no_children})$ and let μ be the added formula $\mu = \neg(\text{has_no_children})$. The only model of ψ is

$$I = \{\text{young}, \text{single}, \text{has_no_children}\}.$$

⁴ $\min$ is here the minimal numerical value of the distance between two interpretations.

The models of μ are:

$$\begin{aligned} J_1 &= \{\text{young}, \text{single}, \neg(\text{has_no_children})\}, \\ J_2 &= \{\text{young}, \neg(\text{single}), \neg(\text{has_no_children})\}, \\ J_3 &= \{\neg(\text{young}), \text{single}, \neg(\text{has_no_children})\} \text{ and} \\ J_4 &= \{\neg(\text{young}), \neg(\text{single}), \neg(\text{has_no_children})\}. \end{aligned}$$

Using the sets of propositional variables on which two interpretations differ,

$$\begin{aligned} \text{diff}(I, J_1) &= \{\text{has_no_children}\}, \\ \text{diff}(I, J_2) &= \{\text{single}, \text{has_no_children}\}, \\ \text{diff}(I, J_3) &= \{\text{young}, \text{has_no_children}\}, \\ \text{diff}(I, J_4) &= \{\text{young}, \text{single}, \text{has_no_children}\}, \end{aligned}$$

and

$$\text{diff}(I, J_2) \subset \text{diff}(I, J_4), \text{diff}(I, J_3) \subset \text{diff}(I, J_4), \text{diff}(I, J_1) \subset \text{diff}(I, J_2), \text{diff}(I, J_1) \subset \text{diff}(I, J_3),$$

therefore $J_1 = \{\text{young}, \text{single}, \neg(\text{has_no_children})\}$ is the model of $\psi \circ_B \mu$.

In Borgida's revision operation, the principle of minimal change takes the form of the minimal subset of propositional variables which change. Since the added formula is $\neg(\text{has_no_children})$, and the model of ψ is $I = \{\text{young}, \text{single}, \text{has_no_children}\}$, the model of μ which only differs by one propositional variable is $J_1 = \{\text{young}, \text{single}, \neg(\text{has_no_children})\}$.

5.2.3 Revision according to Winslett

The revision operation proposed by Winslett, the so-called possible world approach, is defined within the context of reasoning about actions. The same metric as that of Borgida is used, but the definition of the revision operation differs because the consistent and inconsistent cases do not operate the same way. More formally:

$$\text{Mod}(\psi \circ_{\text{pma}} \mu) = \text{Min}_{I \in \text{Mod}(\psi)}(\text{diff}(I, \mu)).$$

This operation satisfies (R1) and the postulates (R3)–(R5). The postulate (R2) is not satisfied; when $\psi \wedge \mu$ is consistent, $\psi \circ_{\text{pma}} \mu$ may not be equivalent to $\psi \wedge \mu$. This is desirable within the framework of reasoning about actions. The non-satisfaction of this postulate initiated a work which allowed, a posteriori, distinguishing between revision and update operations.

Example 13 The previous example is not suitable to illustrate Winslett's approach within the context of actions. So let us consider a very simple and famous example, in which the knowledge base specifies that in a room the window and the door may not be both open, d_open represents the fact that the door is open and w_open the fact that the window is open. The knowledge base is represented by the formula

$$\psi = (d_open \wedge \neg(w_open)) \vee (\neg(d_open) \wedge w_open).$$

Suppose now that this knowledge base is revised by the action of somebody opening the door represented by the formula $\mu = d_open$.

The models of ψ are $I_1 = \{d_open, \neg(w_open)\}$ and $I_2 = \{\neg(d_open), w_open\}$.

The models of μ are $J_1 = \{d_open, w_open\}$ and $J_2 = \{d_open, \neg(w_open)\}$.

Borgida's revision operation yields

$$\psi \circ_B \mu = \psi \wedge d_open$$

and $J_2 = \{d_open, \neg(w_open)\}$ is the model of $\psi \circ_B \mu$. This means that when somebody opens the door, the window is automatically closed, which is not desirable within the context of actions.

On the other hand, Winslett's revision operation finds for each model of ψ the closest models of μ , $\text{diff}(I_1, J_1) = \emptyset$, $\text{diff}(I_1, J_2) = \{w_open\}$, $\text{diff}(I_2, J_1) = \{d_open, w_open\}$, $\text{diff}(I_2, J_2) = \{d_open\}$, therefore $J_1 = \{d_open, w_open\}$ and $J_2 = \{d_open, \neg(w_open)\}$ are the two models of $\psi \circ_{\text{pma}} \mu$. This means that when somebody opens the door, the window may or may not be closed.

5.2.4 Integrity constraints

Though stemming from the principle of irrelevance of syntax, semantic approaches must keep integrity constraints, which specify that some formulae, for instance physical laws, must always be satisfied. Let IC be a formula representing the set of these integrity constraints, the revision operation taking those into account is defined by:

$$\psi \circ^{IC} \mu \equiv \psi \circ (\mu \wedge IC).$$

Semantic approaches are of a great interest from a theoretical point of view; however, the *KM* postulates as well as the *AGM* postulates are non-constructive since they do not provide a method for the construction of revision operations. On the other hand, in most cases it is generally impossible to explicitly represent the revised base. This seems important when one has to deal with iterated revision. The satisfaction of the principle of minimal change involves defining orders either on formulae or on interpretations. These orders are external to the language and remain relatively arbitrary. It would be preferable to define an order involved by the revision process itself. In the literature, works on revision generally operate within the context of consistency preservation. Nevertheless, reality is more complex and one often needs the representation of temporarily inconsistent situations. Semantic approaches stem from the principle of the irrelevance of syntax; notably two equivalent knowledge bases are revised in the same way which in many examples (such as the famous restaurant example proposed by Hansson) contradicts common-sense reasoning. It can be noted that this approach still takes syntax into account by the specific treatment of integrity constraints which specify that some formulae cannot be removed.

5.3 Spohn's Ordinal Conditional Function

Spohn (1987) provides a function, called the *ordinal conditional function*, which associates with each interpretation of a propositional formula an ordinal number. This function allows the expression of degrees of plausibility; the smaller the ordinal, the more plausible the interpretation. Within the context of revision, the function either decreases or increases the ordinal number corresponding to an interpretation, according to whether it either is or is not a model of the added formula, respectively.

5.3.1 Ranking

More formally, an ordinal conditional function, also called a ranking, denoted κ , is a function from a given set of interpretations into the class of ordinals, such that some interpretations are assigned the smallest ordinal, 0.

The ranking is extended from interpretations to propositional formulae: it is the smallest rank assigned to an interpretation that satisfies the formula. More formally, let μ be a formula and I be an interpretation, $\kappa(\mu) = \min_{I \in \text{Mod}(\mu)} \kappa(I)$, consequently, $\kappa(\mu \vee \nu) = \min(\kappa(\mu), \kappa(\nu))$.

A ranking accepts a formula μ if its negation is implausible, i.e. $\kappa(\neg\mu) > 0$. The set of formulae accepted by the ranking is characterised by the set of models assigned the ordinal 0.

5.3.2 Revision according to Spohn

The knowledge base is represented by a set of ranked models. The added information is represented by a pair (μ, m) , where μ is a propositional formula and m is the post-revision degree of plausibility of μ . Revising the knowledge base by a new item of information, also called (μ, m) -conditionalisation of κ , occurs, changing the ranking of interpretations. For all I in $\mathcal{I}$,

$$\kappa_{(\mu, m)}(I) = \begin{cases} \kappa(I) - \kappa(\mu), & \text{if } I \in \text{Mod}(\mu); \\ \kappa(I) - \kappa(\neg\mu) + m, & \text{if } I \notin \text{Mod}(\mu). \end{cases}$$

This revision operation satisfies the postulates (R1)–(R6).

Example 14 Let ψ be the unique formula representing the knowledge base $\psi = \text{young} \wedge \text{single} \wedge (\text{young} \wedge \text{single} \rightarrow \text{has_no_children})$, the only model of ψ is $I = \{\text{young}, \text{single}, \text{has_no_children}\}$ and the corresponding ranking is $\kappa(I) = 0$ and $\forall I' \in \mathcal{I}, I' \neq I, \kappa(I') = 1$.

The revision of ψ by the formula $\mu = \neg(\text{has_no_children})$ with post-revision degree of plausibility 3 leads to the modification of the previous ranking. The models of μ are:

- $J_1 = \{\text{young}, \text{single}, \neg(\text{has_no_children})\},$
- $J_2 = \{\text{young}, \neg(\text{single}), \neg(\text{has_no_children})\},$
- $J_3 = \{\neg(\text{young}), \text{single}, \neg(\text{has_no_children})\}$ and
- $J_4 = \{\neg(\text{young}), \neg(\text{single}), \neg(\text{has_no_children})\}.$

The modified ranking is $\kappa_{(\mu, 3)}(I) = 2, \kappa_{(\mu, 3)}(J_1) = 0, \kappa_{(\mu, 3)}(J_2) = 0, \kappa_{(\mu, 3)}(J_3) = 0, \kappa_{(\mu, 3)}(J_4) = 0$, and $\forall I' \in \mathcal{I}, I' \neq I, I' \neq J_i, i \in \{1, 4\} \kappa_{(\mu, 3)}(I') = 3$.

Consequently,

- $J_1 = \{\text{young}, \text{single}, \neg(\text{has_no_children})\},$
- $J_2 = \{\text{young}, \neg(\text{single}), \neg(\text{has_no_children})\},$
- $J_3 = \{\neg(\text{young}), \text{single}, \neg(\text{has_no_children})\}$ and
- $J_4 = \{\neg(\text{young}), \neg(\text{single}), \neg(\text{has_no_children})\}$

are the models of ψ revised by μ with post-revision degree of plausibility 3.

The result of revision gives the four models of the added formula, because these interpretations have the same ranking at the previous step. The principle of minimal change takes the form of the minimal change of the previous ranking, and Spohn's approach to revision is suitable for iterated revision.

6 Mixed approaches

6.1 Mixed approach according to Willard and Li

A compromise between syntactic and semantic approaches is proposed by Willard and Li (1990). The formulae are prioritised such that explicit formulae have higher priority than derived formulae, and among explicit formulae, complex formulae or rules have higher priority than atomic formulae or facts.

This method draws inspiration from that of Nebel. However, the clausal form is used in order to define a normal form of the knowledge base. This normal form involves both the standard clausal form and some explicitly given *underlined* formulae (Levy, 1994).

6.1.1 Preliminaries

In this approach the knowledge base is a set of clauses denoted B .

Definition 1 Let μ and ν be two clauses, μ subsumes ν iff any literal of ν is a literal of μ .

Any set of clauses has a unique subset of minimal clauses for subsumption. $Sb(B)$ denotes the set of minimal clauses of $Cons(B)$, more formally: $Sb(B) = \{\mu, \text{ such that } B \vdash \mu \text{ and } \neg \exists \nu \text{ such that } B \vdash \nu \text{ and } \nu \text{ subsumes } \mu \text{ and } \mu \text{ does not subsume } \nu\}$.

6.1.2 Survey of the method

An outline of the method is given below, for more details see Willard & Li (1990) and Levy (1994). The partial closure of B , denoted $Pc(B)$ and defined by $Pc(B) = Sb(B) \cup B$, is constructed such that two equivalent bases only differ by the explicitly given set of clauses.

A clause is underlined if it does not belong to $Sb(B)$ and all the underlined clauses of $Pc(B)$ belongs to B .

A base is *partially closed* iff $Pc(B) = B$, in other words $Sb(B) \subseteq B$.

In order to construct a normal form of the knowledge base, the notion of *extended intersection* between bases is then introduced such that the extended intersection of any finite number of partially

closed bases is a partially closed base. More formally, let $B_i, 1 \leq i \leq n$ be a set of bases, $\hat{\cap}_{1 \leq i \leq n} B_i$ denotes the extended intersection between these bases:

$$\hat{\cap}_{1 \leq i \leq n} B_i = (\cap_{1 \leq i \leq n} B_i) \cup (\{\vee_{1 \leq i \leq n} \mu_i, \text{ with } \mu_i \in (B_i - (\cap_{1 \leq i \leq n} B_i))\}).$$

Subsumption is extended to underlined formulae as follows: let B_i and B_j be two bases, $B_i \leq B_j$ iff $B_j \vdash B_i$ and B_j contains all the underlined formulae of B_i .

The AGM postulates are then reformulated within the context of partially closed bases. Let B be a partially closed base, let A and C be two sets of clauses, $B \star A$ denotes the partially closed base representing the base B revised by A and $B + A$ denotes the partial closure of $B \cup A$

- (R★1) $B \star A$ is a partially closed theory.
- (R★2) $A \leq B \star A$.
- (R★3) $B \star A \leq B + A$.
- (R★4) $B + A \leq B \star A$ if $B + A$ is consistent.
- (R★5) $B \star A$ is inconsistent if and only if A is.
- (R★6) If $Pc(A) = Pc(C)$ then $B \star A = B \star C$.
- (R★7) $B \star (A \cup C) \leq (B \star A) + C$.
- (R★8) If $(B \star A) + C$ is inconsistent then $(B \star A) + C \leq B \star (A \cap B)$.

The proposed revision operation satisfying the above postulates follows several steps. It first defines a normalising operation of the knowledge base which involves both a standard clausal form and some explicitly given underlined formulae (Levy, 1994).

$$B^{nu} A = Sb(B) \cup \{ \mu, B \vdash \mu \text{ and } \mu \in A \}.$$

The deductive closure of the partially closed base $B^{nu} A$ is $Cons(B)$. In $Cons(B)$ clauses of $A \cap Cons(B)$ are underlined.

$MS(B, A)$ is defined as the set of maximal consistent subsets containing A . A base, called $Sum(B, A)$, and such that both A and $\cap_{B \in MS(B, A)} Cons(B)$ are implied, is then constructed.

$$Sum(B, A) = (\hat{\cap}(MS(B, A)))^{nu}(B \cup A).$$

Finally the revision operation is defined by:

$$B \star A = Sum(B, A).$$

This revision operation is the analogue to the full meet revision defined by Gärdenfors. Another revision operation analogous to the partial meet revision is proposed by the authors. It is defined from a selection function from preferences between positive and negative literals chosen in the set $MS(B, A)$. For more details see Willard & Li (1990) and Levy (1994).

Example 15 In order to illustrate this approach, let B be a knowledge base expressed by a standard normal clause, $B = \{young, single, \neg(young) \vee (\neg(single) \vee has_no_children)\}$ and A be the added clause, $A = \{\neg(has_no_children)\}$.

Let $MS(B, \{\neg(has_no_children)\})$ be the set of maximal consistent subsets with $\{\neg(has_no_children)\}$ then $MS(B, \{\neg(has_no_children)\}) = \{\{young, single\}, \{young, \neg(young) \vee (\neg(single) \vee has_no_children)\}, \{b, \neg(young) \vee (\neg(single) \vee has_no_children)\}\}$.

Remember that the revision operation is defined by:

$$Sum(B, \{\neg(has_no_children)\}) = (\hat{\cap} (MS(B, \{\neg(has_no_children)\})))^{nu}(B \cup \{\neg(has_no_children)\}).$$

As the extended intersection is defined by:

$$\hat{\cap}_{1 \leq i \leq n} B_i = (\cup_{1 \leq i \leq n} B_i) \cup (\{\vee_{1 \leq i \leq n} \mu_i, \text{ with } \mu_i \in (B_i - (\cap_{1 \leq i \leq n} B_i))\})$$

the extended intersection of

$MS(B, \{\neg(\text{has_no_children})\})$ thus equals

$$\emptyset \cup ((\text{young} \wedge \text{single}) \vee (\text{young} \wedge (\neg(\text{young}) \vee (\neg(\text{single}) \vee \text{has_no_children}))) \vee (\text{single} \wedge (\neg(\text{young}) \vee (\neg(\text{single}) \vee (\text{has_no_children}))))))$$

therefore,

$$\hat{\cap}(MS(B, \{\neg(\text{has_no_children})\})) = \{\text{young} \vee \text{single}\}.$$

Since

$$\hat{\cap}(MS(B, \{\neg(\text{has_no_children})\}))^{nu} = Sb(MS(B, \{\neg(\text{has_no_children})\}) \cup \{\mu, MS(B, \{\neg(\text{has_no_children})\}) \vdash \mu, \text{ and } \mu \in \{\neg(\text{has_no_children})\},$$

therefore

$$Sb(MS(B, \{\neg(\text{has_no_children})\})) = \{\text{young} \vee \text{single}\} \text{ and } \{\mu, MS(B, \{\neg(\text{has_no_children})\}) \vdash \mu, \text{ and } \mu \in \{\neg(\text{has_no_children})\} = \{\neg(\text{has_no_children})\}.$$

Hence:

$$\begin{aligned} B \star \{\neg(\text{has_no_children})\} &= Sum(B, \{\neg(\text{has_no_children})\}) \\ &= \{\text{young} \vee \text{single}, \neg(\text{has_no_children})\}. \end{aligned}$$

Revising by $\neg(\text{has_no_children})$ we cannot any more have both young and single in the revised knowledge base. This revision operation keeps $\text{young} \vee \text{single}$ which makes sense and follows the principle of minimal change.

6.2 Revision in extended propositional calculus

The principle of irrelevance of syntax, from which semantic approaches stem, fails in several examples of common-sense reasoning. This is one of the reasons why mixed methods have been developed. On the other hand, it seems important to represent explicitly the revised knowledge base, particularly within the context of iterated revision.

Following this point of view, Papini and Rauzy (1995, see also Papini, 1996a) proposed to extend the propositional calculus with new connectives in order to define revision operations. This approach is a mixed one in the sense that it can be placed within the framework of semantic approaches, because revision operations involve models, but meanwhile syntax is used to express preferences between models.

Propositional calculus is extended with two new modalities, $\flat$ and $\sharp$, which make it possible to represent revised knowledge bases, especially in the case of iterated revision. The unary $\flat$ and $\sharp$ operators respectively weaken or strengthen the formula they prefix. The $\flat$ operator permits one to define a revision operation $((\psi \circ_{\flat} \mu = (\flat \psi) \wedge \mu))$ when the new item of information is preferred, e.g. in case of social number or meteorological data. Nevertheless, in certain applications, the initial knowledge base has to be temporarily preferred with respect to the new item of information. This is the reason why a dual operator, called $\sharp$, is defined; it permits one to define another revision operation $((\psi \circ_{\sharp} \mu = (\sharp \psi) \wedge \mu))$.

As within the framework of semantic approaches, a knowledge base is represented by a unique formula ψ . The semantics of the extended propositional calculus are defined by means of weighted interpretations. Any interpretation of the formula is linked with a weight. The weight of a formula, for a given interpretation, is a polynomial of $\mathbb{N}[-x, x]$. A formula satisfied by an interpretation has the null polynomial as weight. On the other hand, the weight of an unsatisfied formula is a polynomial whose constant coefficient is non-zero. The weights are ordered according to a total order induced by the revision process itself, taking into account both the syntax and the history of the knowledge base. The preferred non-classical models of the formula representing the revised knowledge base are the interpretations assigned the smallest weight.

Example 16 Let ψ be the unique formula representing the knowledge base $\psi = \text{young} \wedge \text{single} \wedge (\neg(\text{young}) \vee \neg(\text{single}) \vee \text{has_no_children})$ and $\mu = \neg(\text{has_no_children})$. Assuming that the more recent an item of information the more certain it is, the revised knowledge base is represented by:

$$\psi \circ_b \mu = b(\text{young} \wedge \text{single} \wedge (\neg(\text{young}) \vee \neg(\text{single}) \vee \text{has_no_children})) \wedge \neg(\text{has_no_children}).$$

Let the different interpretations be:

- σ_1 : [$\text{young} \leftarrow 0, \text{single} \leftarrow 0, \text{has_no_children} \leftarrow 0$],
- σ_2 : [$\text{young} \leftarrow 0, \text{single} \leftarrow 0, \text{has_no_children} \leftarrow 1$],
- σ_3 : [$\text{young} \leftarrow 0, \text{single} \leftarrow 1, \text{has_no_children} \leftarrow 0$],
- σ_4 : [$\text{young} \leftarrow 0, \text{single} \leftarrow 1, \text{has_no_children} \leftarrow 1$],
- σ_5 : [$\text{young} \leftarrow 1, \text{single} \leftarrow 0, \text{has_no_children} \leftarrow 0$],
- σ_6 : [$\text{young} \leftarrow 1, \text{single} \leftarrow 0, \text{has_no_children} \leftarrow 1$],
- σ_7 : [$\text{young} \leftarrow 1, \text{single} \leftarrow 1, \text{has_no_children} \leftarrow 0$],
- σ_8 : [$\text{young} \leftarrow 1, \text{single} \leftarrow 1, \text{has_no_children} \leftarrow 1$].

Let ψ be a formula of the extended propositional calculus, a weight, denoted $w_{\sigma_i, \psi}(-x, x)$, is computed for any interpretation σ_i . This involves the notion of polarity, defined in the following way: the polarity of a sub-formula equals 1 if and only if the sub-formula is prefixed by an even number of negations, otherwise it equals 0. If $\sigma_i(\psi) = 1$ then $w_{\sigma_i, \psi}(-x, x) = 0$, because $\text{pol}_\psi(\psi) = 1$, otherwise the weight $w_{\sigma_i, \psi}(-x, x)$ is computed recursively according to the weight of its sub-formulae.

The revised knowledge base is represented by $\psi \circ_b \mu$ and for any interpretation σ_i the corresponding weight, denoted $w_{\sigma_i, \psi \circ_b \mu}(-x, x)$, is computed. The computation of $w_{\sigma_1, \psi \circ_b \mu}(-x, x)$ is now detailed.

σ_1 : [$\text{young} \leftarrow 0, \text{single} \leftarrow 0, \text{has_no_children} \leftarrow 0$], as $\sigma_1(\psi \circ_b \mu) = 0$ and $\text{pol}_{\psi \circ_b \mu}(\psi \circ_b \mu) = 1$. The weight corresponding to σ_1 is computed recursively according to the weight of the sub-formulae of $\psi \circ_b \mu$ as follows:

$$\begin{aligned} w_{\sigma_1, \psi \circ_b \mu}(-x, x) &= w_{\sigma_1, b(\psi)}(-x, x) + w_{\sigma_1, \mu}(-x, x) \\ &= xw_{\sigma_1, \text{young}}(-x, x) + xw_{\sigma_1, \text{single}}(-x, x) \\ &\quad + xw_{\sigma_1, \neg(\text{young}) \vee \neg(\text{single}) \vee \text{has_no_children}}(-x, x) \\ &\quad + w_{\sigma_1, \neg(\text{has_no_children})}(-x, x) \\ &= x + x + 0 + 0 = 2x \end{aligned}$$

The other weights are computed in the same way and:

$$\begin{aligned} w_{\sigma_1, \psi \circ_b \mu}(-x, x) &= x + x + 0 + 0 = 2x \\ w_{\sigma_2, \psi \circ_b \mu}(-x, x) &= x + x + 0 + 1 = 2x + 1 \\ w_{\sigma_3, \psi \circ_b \mu}(-x, x) &= x + 0 + 0 + 0 = x \\ w_{\sigma_4, \psi \circ_b \mu}(-x, x) &= x + 0 + 0 + 1 = x + 1 \\ w_{\sigma_5, \psi \circ_b \mu}(-x, x) &= 0 + x + 0 + 0 = x \\ w_{\sigma_6, \psi \circ_b \mu}(-x, x) &= 0 + x + 0 + 1 = x + 1 \\ w_{\sigma_7, \psi \circ_b \mu}(-x, x) &= 0 + 0 + 0 + x + x + x = 3x \\ w_{\sigma_8, \psi \circ_b \mu}(-x, x) &= 0 + 0 + 0 + 1 = 1 \end{aligned}$$

There are 2 non-classical models whose weights are minimal:

$$\begin{aligned} \sigma_3: [\text{young} \leftarrow 0, \text{single} \leftarrow 1, \text{has_no_children} \leftarrow 0] \text{ with } w_{\sigma_3, \psi \circ_b \mu}(-x, x) &= x \\ \sigma_5: [\text{young} \leftarrow 1, \text{single} \leftarrow 0, \text{has_no_children} \leftarrow 0] \text{ with } w_{\sigma_5, \psi \circ_b \mu}(-x, x) &= x \end{aligned}$$

Hence σ_3 and σ_5 are the two preferred non-classical models of $\psi \circ_b \mu$.

The principle of minimal change is followed by the fact that the models of the revised formula differ from σ_8 , the only model of the initial knowledge base, by either the propositional variable *young* or the propositional variable *single*.

7 Revision and update

Katsuno and Mendelzon (1991b) throw light on the difference between revision and update operations, while unifying in a common framework semantic approaches of revision. Revision occurs when a new item of information is introduced into a static world; in contrast update occurs when the world itself evolves. Revising a set of possible worlds, representing the knowledge, comes to make this set of worlds evolve as a whole towards the closest set of possible worlds satisfying the new item of information. On the other hand, updating a set of possible worlds, representing the knowledge, comes to make each possible world evolve locally towards the closest world satisfying the new item of information (Sur le Collectif, 1995).

Example 17 *Let us recall a very simple and famous example within the context of actions. The knowledge base, represented by the formula*

$$\psi = (d_open \wedge \neg w_open) \vee (\neg d_open \wedge w_open),$$

specifies that in a room the window and the door may not be both open, where d_open represents the fact that the door is open and w_open the fact that the window is open. The models of ψ are $I_1 = \{d_open, \neg w_open\}$ and $I_2 = \{\neg d_open, w_open\}$. The situation described is not a static one, it can evolve as a result of an agent's action. In updating, the situation itself is changing, as opposed to revision, where the agent's perception of a situation is changing. Suppose now that somebody opens the door, represented by the formula $\mu = d_open$. The models of μ are $J_1 = \{d_open, w_open\}$ and $J_2 = \{d_open, \neg(w_open)\}$. Two cases are now considered, according to whether the window was either open or not open before the execution of the action. In each case, the closest models of μ are determined, and the final result of the update consists of the set of models obtained in both cases. More precisely, the model $J_1 = \{d_open, w_open\}$ of ψ corresponds to the case where the window was open before the execution of the action and the model $J_2 = \{d_open, \neg(w_open)\}$ of ψ corresponds to the case where the window was closed before the execution of the action. Using Borgida's distance, denoted $diff$,⁵ to measure the proximity between models,

$$\begin{aligned} diff(I_1, J_1) &= \emptyset, \\ diff(I_1, J_2) &= \{w_open\}, \\ diff(I_2, J_1) &= \{d_open, w_open\}, \\ diff(I_2, J_2) &= \{d_open\}, \end{aligned}$$

therefore

$J_1 = \{d_open, w_open\}$ and $J_2 = \{d_open, \neg(w_open)\}$ are the two models of $\psi \diamond \mu$. This means that when somebody opens the door, the window may or may not be closed.

7.1 KM postulates

As in the revision case, Katsuno and Mendelzon (1991a) formulated postulates, known as KM postulates, that an updated knowledge base has to satisfy.

KM postulates Let ψ , ϕ and μ be formulae,

- (U1) $\psi \diamond \mu$ implies μ .
- (U2) If ψ implies μ , then $\psi \diamond \mu \equiv \psi$.
- (U3) If μ and ψ are satisfiable, then so is $\psi \diamond \mu$.
- (U4) If $\psi_1 \equiv \psi_2$ and $\mu_1 \equiv \mu_2$,
then, $\psi_1 \diamond \mu_1 \equiv \psi_2 \diamond \mu_2$.
- (U5) $(\psi \diamond \mu) \wedge \phi$ implies $\psi \diamond (\mu \wedge \phi)$.

⁵The distance $diff$ is defined in section 5.

- (U6) If $\psi \diamond \mu_1$ implies μ_2 and $\psi \diamond \mu_2$ implies μ_1 ,
 then $\psi \diamond \mu_1 \equiv \psi \diamond \mu_2$.
 (U7) If ψ is complete, then
 $(\psi \diamond \mu_1) \wedge (\psi \diamond \mu_2)$ implies $\psi \diamond (\mu_1 \vee \mu_2)$.
 (U8) If $(\psi_1 \wedge \psi_2) \diamond \mu \equiv (\psi_1 \diamond \mu) \wedge (\psi_2 \diamond \mu)$.

The postulates (U1)–(U5) directly correspond to postulates stated for revision; the postulate (U2) means that if μ may be derived from ψ , then updating by μ does not modify ψ . However, if ψ is consistent (U2) is weaker than (R2). (U6) specifies that each possible world is equally considered. (U7) suggests that if the initial knowledge base contains no uncertain information, then if a possible world results both from updating by μ_1 and from updating by μ_2 it also results from updating by $\mu_1 \vee \mu_2$. (U8) illustrates the local evolution of each possible world towards the closest worlds satisfying μ .

7.2 Principle of minimal change

The principle of minimal change leads to the definition of orders between interpretations.

Let I be the set of all the interpretations and $Mod(\psi)$ be the set of models of ψ . In contrast with the revision case, in which a pre-order is associated with ψ , a partial pre-order on I , denoted $\leq_I$, is associated with each interpretation I . The relation $<_I$ is defined from $\leq_I$ and the pre-order is faithful if it verifies the following condition:

$$\forall J, J \in I, \text{ if } I \neq J \text{ then } <_I J.$$

Let ψ be a formula representing a knowledge base, Katsuno and Mendelzon showed that an updating operation satisfies the postulates (U1)–(U8) if and only if there exists a partial pre-order $\leq_I$ such that

$$Mod(\psi \diamond \mu) = \bigcup_{I \in Mod(\psi)} Min(Mod(\mu), \leq_I).$$

8 Numerical approaches

Change operations may also be considered within the framework of probability or possibility theories. These numerical formalisms (Sombe, 1992) are used in order to represent uncertainty. They describe knowledge states in terms of possible states of the world. These worlds, denoted $\omega \in \Omega$, are mutually exclusive and correspond to models in logical formalisms.

Both in probability theory and in possibility theory, to each world ω is attached a degree, denoted $d(\omega)$, which estimates the extent to which ω may represent the real state of the world. The distribution d is such that $d(\omega) \in [0, 1]$. $d(\omega) = 0$ means that the possible world ω does not correspond to a real state of the world. In probability theory, $d(\omega) = 1$ means that ω is the real state of the world. However, in possibility theory, $d(\omega) = 1$ expresses that nothing prevents ω from being the real world. Within the framework of possibility, ω certainly is the real state of the world if $\forall \omega' \neq \omega, d(\omega') = 0$, i.e. information not induced by $d(\omega) = 1$, although within the framework of probability theory $d(\omega) = 1 \rightarrow d(\omega') = 0, \forall \omega' \neq \omega$.

In numerical formalisms, the change of a state of knowledge by the introduction of new information, stating that the real world is in $A \subseteq \Omega$, corresponds to a modification of the distribution going from d to d' .

Generally, the distribution d' results from a conditioning operation $d'(\omega) = d(\omega | A)$. The change operation has to keep three principles: (1) d' is of the same nature as d (preservation of the representation principles), (2) what is observed is held as certain after revision, i.e. $\forall \omega \notin A, d'(\omega) = 0$, (3) principle of minimal change, i.e. the distance between d and d' has to be minimised. These principles are analogous to the principles characterising the revision by Gärdenfors, Alchourron and Makinson.

8.1 Bayesian revision

In this approach (Pearl, 1988), a state of knowledge is characterised by a probability distribution p , such that $\sum_{\omega \in \Omega} p(\omega) = 1$. A probability measure is attached to the probability distribution $\forall E \subseteq \Omega$, $P(E) = \sum_{\omega \in E} p(\omega)$. If $p(\omega) = 1$ then the state of the world certainly is ω , hence $\forall \omega' \neq \omega$, $p(\omega') = 0$. The Bayes formula $P(B | A) = \frac{P(A \cap B)}{P(A)}$ permits one to modify the probability distribution after the arrival of a certain new item of information. The conditioning of p for the observation A is then:

$$p(\omega | A) = \frac{p(\omega)}{P(A)}, \text{ if } \omega \in A; \quad (4)$$

$$= 0, \text{ if } \omega \notin A. \quad (5)$$

Example 18 Let K be the knowledge base $K = \{\text{young}, \text{single}, (\text{young} \wedge \text{single} \rightarrow \text{has_no_children})\}$ and let $\neg(\text{has_no_children})$ be the added information. The possible states of the world are:

$$\begin{aligned} \omega_1 &= (\neg(\text{young}) \wedge \neg(\text{single}) \wedge \neg(\text{has_no_children})) \\ \omega_2 &= (\neg(\text{young}) \wedge \neg(\text{single}) \wedge \text{has_no_children}) \\ \omega_3 &= (\neg(\text{young}) \wedge \text{single} \wedge \neg(\text{has_no_children})) \\ \omega_4 &= (\neg(\text{young}) \wedge \text{single} \wedge \text{has_no_children}) \\ \omega_5 &= (\text{young} \wedge \neg(\text{single}) \wedge \neg(\text{has_no_children})) \\ \omega_6 &= (\text{young} \wedge \neg(\text{single}) \wedge \text{has_no_children}) \\ \omega_7 &= (\text{young} \wedge \text{single} \wedge \neg(\text{has_no_children})) \\ \omega_8 &= (\text{young} \wedge \text{single} \wedge \text{has_no_children}) \end{aligned}$$

and their corresponding probabilities are:

$$\begin{aligned} \omega_1 &= p(\omega_1) = 0, \\ \omega_2 &= p(\omega_2) = 0.1, \\ \omega_3 &= p(\omega_3) = 0.2, \\ \omega_4 &= p(\omega_4) = 0.1, \\ \omega_5 &= p(\omega_5) = 0.2, \\ \omega_6 &= p(\omega_6) = 0.3, \\ \omega_7 &= p(\omega_7) = 0.1, \\ \omega_8 &= p(\omega_8) = 0. \end{aligned}$$

Let A be the set of worlds in which $\neg(\text{has_no_children})$ is certain, $A = \{\omega_1, \omega_3, \omega_5, \omega_7\}$. Upon the occurrence of A , the Bayes formula yields:

$$\begin{aligned} p(\omega_1 | A) &= \frac{p(\omega_1)}{p(A)} = 0, & p(\omega_3 | A) &= \frac{p(\omega_3)}{p(A)} = \frac{2}{5}, \\ p(\omega_5 | A) &= \frac{p(\omega_5)}{p(A)} = \frac{2}{5}, & p(\omega_7 | A) &= \frac{p(\omega_7)}{p(A)} = \frac{1}{5}. \end{aligned}$$

Hence the most probable worlds are ω_3 and ω_5 .

The principle of minimal change is followed by the fact that the models of the revised formula differ from ω_8 , the only model of the initial knowledge base, by either the propositional variable *young* or the propositional variable *single*.

8.2 Revision by imaging

Lewis (1976) proposed an imaging rule which moves the probability $p(\omega)$ from each world ω to the closest world ω_A in A . The image p_A of p in A is obtained as follows:

$$\forall \omega' \in A, p_A(\omega') = \sum_{\omega=\omega_A} p(\omega), \quad (6)$$

$$\forall \omega' \notin A, p_A(\omega') = 0, \quad (7)$$

where ω_A is the closest world ω in A . If there are several closest worlds to ω in A , the above formula can be generalised sharing $p(\omega)$ among the various worlds $\omega' \in A(\omega)$.

Example 19 Let K be the knowledge base $K = \{\text{young}, \text{single}, (\text{young} \wedge \text{single} \rightarrow \text{has_no_children})\}$ and let $\neg(\text{has_no_children})$ be the added information. The possible states of the world are:

- ω_1 $(\neg(\text{young}) \wedge \neg(\text{single}) \wedge \neg(\text{has_no_children}))$,
- ω_2 $(\neg(\text{young}) \wedge \neg(\text{single}) \wedge \text{has_no_children})$,
- ω_3 $(\neg(\text{young}) \wedge \text{single} \wedge \neg(\text{has_no_children}))$,
- ω_4 $(\neg(\text{young}) \wedge \text{single} \wedge \text{has_no_children})$,
- ω_5 $(\text{young} \wedge \neg(\text{single}) \wedge \neg(\text{has_no_children}))$,
- ω_6 $(\text{young} \wedge \neg(\text{single}) \wedge \text{has_no_children})$,
- ω_7 $(\text{young} \wedge \text{single} \wedge \neg(\text{has_no_children}))$,
- ω_8 $(\text{young} \wedge \text{single} \wedge \text{has_no_children})$.

and their corresponding probabilities are:

- ω_1 $p(\omega_1) = 0$,
- ω_2 $p(\omega_2) = 0.1$,
- ω_3 $p(\omega_3) = 0.2$,
- ω_4 $p(\omega_4) = 0.1$,
- ω_5 $p(\omega_5) = 0.2$,
- ω_6 $p(\omega_6) = 0.3$,
- ω_7 $p(\omega_7) = 0.1$,
- ω_8 $p(\omega_8) = 0$.

Let A be the set of worlds in which $\neg(\text{has_no_children})$ is certain, $A = \{\omega_1, \omega_3, \omega_5, \omega_7\}$. Revision by imaging moves $p(\omega)$ from each world ω to the closest world ω_A in A . The closeness between worlds is expressed by the Hamming distance. The closest world of ω_1 is ω_2 thus $p_A(\omega_1) = p(\omega_2) = 0.1$, the closest worlds of ω_3 are $\omega_2, \omega_4, \omega_8$ thus $p_A(\omega_3) = p(\omega_2) + p(\omega_4) + p(\omega_8) = 0.2$, the closest world of ω_5 is ω_6 thus $p_A(\omega_5) = p(\omega_6) = 0.2$, the closest world of ω_7 is ω_8 thus $p_A(\omega_7) = p(\omega_8) = 0$. Hence the most probable worlds are ω_3 and ω_5 .

Like in the previous revision operation, the principle of minimal change is followed by the fact that the models of the revised formula differ from ω_8 , the only model of the initial knowledge base, by either the propositional variable *young* or the propositional variable *single*.

The analogy between revision methods based on AGM postulates and KM postulates, on one hand, and the Bayesian revision and revision by imaging, on the other hand, can be noted.

These two kinds of revision can also be defined within a possibilistic framework Dubois *et al.*, (1994). Possibilistic Bayesian revision exactly satisfies the counterpart of AGM postulates and possibilistic revision by imaging satisfies the counterpart of KM postulates for updating. This is not the case of the probabilistic approach which only satisfies the counterpart of some postulates.

9 Revision and non-monotony

Non-monotonic reasoning formalisms introduced for solving general problems are closely con-

nected with revision. Makinson and Gärdenfors (1991) established a precise relationship between these two theories.

On the one hand, a non-monotonic inference relation can be defined from the revision theory by considering the revision of a deductively closed theory T by a formula μ in order to hold the theory $T \star \mu$. This can be viewed as a certain form of non-monotonic inference starting from μ in the context of T . On the other hand, the non-monotonic inference of q from p can be considered as the fact that q belongs to the result of the revision of some theory T by p .

Let $\sim$ be the non-monotonic inference

$$p \sim_T q \leftrightarrow q \in T \star p,$$

the AGM postulates are translated as follows, where $\top$ denotes the tautology:

- (NM $\star$ 1) $Th_{NM}(p) = Th(Th_{NM}(p))$.
- (NM $\star$ 2) $p \sim p$.
- (NM $\star$ 3) $p \sim q \rightarrow \top \sim (p \sim q)$.
- (NM $\star$ 4) If $\top \sim \neg p$ and $\top \sim p \rightarrow q$ then $p \sim q$.
- (NM $\star$ 7) If $p \wedge q \sim r$ then $q \sim p \rightarrow r$.
- (NM $\star$ 8) If $q \sim \neg p$ and $q \sim p \rightarrow r$ then $p \wedge q \sim r$.

Referring to the original AGM postulates, $T \star A$ is a theory becomes (NM $\star$ 1). $A \in T \star A$ becomes (NM $\star$ 2); $T \star A \subseteq T + A$ becomes (NM $\star$ 3); If $\neg A \notin T$ then $T \star A = T + A$ becomes (NM $\star$ 4); $T \star (A \wedge B) \subseteq (T \star A) + B$ becomes (NM $\star$ 7); If $\sim \neg B \notin T \star A$ then $(T \star A) + B = T \star (A \wedge B)$ becomes (NM $\star$ 8).

Not all non-monotonic inferences satisfy these postulates; Reiter's default theory, for instance, does not satisfy the third postulate.

10 Revision and constraints

Revision and constraints can be linked via dynamic CSPs.⁶ The CSP formalism introduced by Montanari (1974) permits the representation and the treatment of many problems in artificial intelligence, within a static environment. Nevertheless, this approach is not suitable for dealing with dynamic problems. This is the reason why the notion of dynamic CSP (DCSP) has been introduced by Dechter and Dechter (1988). A dynamic CSP is a sequence of CSPs, in which each element differs from the previous by the addition or deletion of a constraint.

In order to provide efficient and stable solutions, the adopted process consists in storing the results and using them again within the framework of close problems (Sur le Collectif, 1995). Each deletion of constraint preserves the solutions and each addition of constraint preserves the deductions. However, an addition of constraint may reject a solution and a deletion of constraint may disprove a deduction. Methods allowing one to maintain solutions and/or deductions have to be developed, independently of the nature of change.

The various methods developed in the literature can be divided into three classes (Sur le Collectif, 1995), approaches maintaining consistency, approaches maintaining solutions and approaches using BDDs (Bouquet & Jegou, 1994).⁷

10.1 Approaches maintaining consistency

These stem from filtering methods, in particular filtering by arc-consistency, which permits one to delete from the domains of each variable some values not involved in any solution. Within the framework of dynamic CSPs, in order to avoid filtering again after each modification of the CSP,

⁶Constraint satisfaction problems.

⁷Binary decision diagrams.

the reintroductions of values are restricted. Some methods use (Bessière, 1991) or do not use (Berlandier & Neveu, 1994) justifications for the deletion of values.

10.2 Approaches maintaining solutions

Two families of approaches can be distinguished. For the first one, classical methods of search based on *backtracking* can be improved thanks to production, storage and reuse of justified deduced constraints. The second approach uses repairing methods, allowing one to provide a solution for a CSP from any assignment of the set of variables, more particularly, from a solution provided for a previous CSP (Schiex & Verfaillie, 1994).

10.3 Approaches using BDDs

A BDD (Bryant, 1992) represents a Boolean function by means of a labelled acyclic directed graph. The graph owns a root and two terminal vertices labeled by the truth values 0 and 1 representing the constant Boolean function 0 and 1 respectively. The non-terminal vertices are labelled by propositional variables and the arcs are labeled by the assignments 0 or 1 of the originating variables. Each path from the root to a terminal vertex corresponds to an interpretation of the variables. The truth value of the function corresponding to this interpretation is provided by the label of the terminal vertex.

The BDD approach is used following two main directions. The framework of the BDD is first extended from binary variables to n -ary variables (Vempaty, 1992). The framework of the BDD is then extended from the framework of the CSP to the Boolean framework. The addition of constraint is easily performed, because the construction of a BDD is incremental. The construction of a BDD representing the conjunction of two Boolean functions is performed from the BDD representing each of the Boolean functions. In contrast, the deletion of constraint remains more problematical. The adopted solution in Bouquet & Jegou (1994) consists in preventively treating deletion when the addition of constraint is performed.

11 Overview of recent works on revision

The approaches of revision are numerous; nevertheless none of the proposed revision methods are suitable in every situation. The choice of a method depends mainly on the context of the application within which it is used.

As mentioned in section 7, Katsuno and Mendelzon have shed light on the difference between revision and update operations. They have then characterised update with a set of eight postulates and have provided a representation theorem similar to those defined for revision. Some recent work concentrates on the definition of a unified logical framework in order to either capture both revision and update (Boutilier, 1995) or model how an agent's beliefs change over time (Friedman & Halpern, 1994b).

Furthermore, most of the approaches previously presented in this survey consider one-step revision. However, realistic applications need iteration of the revision process and several recent works have focused on iterated revision.

The AGM paradigm does not support iterated revision, because the underlying preference relation is lost in the process of change, therefore a policy of change is necessary.

Extending Spohn's work, Williams (1994, 1995) refers to the process of changing an underlying preference relation as a transmutation. The added information is a formula α and an ordinal i which represents the information to be accepted with a degree of firmness i . The (α, i) -transmutation of a ranking involves minimal change to the initial ranking such that the formula α is accepted with degree of firmness i . According to the interpretation of the principle of minimum change, governing the policy of change, there result two different forms of transmutations, conditionalisation and adjustment. On the one hand, (α, i) -conditionalisation is based on the relative measure of minimal

change, i.e. all the models of $\neg\alpha$ are shifted such that i is the smallest ordinal corresponding to a model of $\neg\alpha$. The relative ranking between models of α and models of $\neg\alpha$, respectively, is preserved. On the other hand, (α, i) -adjustment is based on the absolute measure of minimal change, i.e. not all the models of $\neg\alpha$ or models of α are shifted. The models of α corresponding to the smallest ordinal are shifted such that their new corresponding ordinal is 0. The models of $\neg\alpha$ corresponding to an ordinal smaller than i are shifted such that their new corresponding ordinal equals i .

Several authors (Friedman & Halpern, 1994a, 1996; Darwich & Pearl, 1997) underline, in the context of iterated revision, the importance of dealing with epistemic states rather than an agent's current beliefs. The epistemic state does not only consist of the agent's current beliefs, but also encodes the strategy that the agent uses to modify his beliefs after learning a new piece of information.

Conditionals of the form *if α were the case then β would be true* represent information that an agent is prepared to adopt, conditioned on future observations. They constitute an important component of the agent's epistemic state. The first connection between revision and conditionals is due to Grove (1988) who interprets a conditional as *revising a knowledge base by α will lead to a state in which β holds*. Consequently, studies of revision and conditional are closely linked. A revision operation determines which conditionals are accepted or rejected; conversely a set of accepted conditionals determines a revision operation.

Several approaches start with the AGM postulates and augment them in order to characterise iterated revision.

In Lehmann's (1995) approach of iterated revision, the agent's epistemic state is a sequence of observations and the revision operation concatenates a new observation to the current epistemic state. This revised framework provides a new set of postulates. Some postulates capture the original AGM postulates, the others are specific to the iteration of the revision process. They stem from the following intuitions: if an agent first learns partial information then the full information, the revision by the partial information cannot influence the result of the revision by the full information; if an agent learns successively the same piece of information, the result of the revision is equivalent to the case where he learns this information once; if an agent successively learns two contradictory pieces of information, he does not retain any of the beliefs obtained during the revision by the intermediate, false information.

Boutilier (1993) takes the agent's epistemic state to consist of a ranking of interpretations. He defines a revision operation, called natural revision, which maps a ranking of interpretations and a formula to a revised ranking of interpretations. This operation tries to minimise changes to ranking and to preserve as many conditionals as the AGM postulate permit. At first glance, this strategy, called *absolute minimization*, seems reasonable, however it can lead to counterintuitive results.

Darwiche and Pearl (1997) show that AGM postulates are too weak to ensure the preservation of conditionals. They reformulate them in terms of epistemic states, and propose additional postulates in order to characterise iterated revision. These postulates rely on the intuition that, first, the relative ordering between the models of the added formula is preserved, then that the relative ordering between the countermodels of the added formula is preserved, and finally that the ranking between models and countermodels of the added formula does not change. As an example of a revision operation satisfying the extended AGM postulates, they slightly modify Spohn's ordinal conditional function.

Friedman and Halpern (1994a) introduce a general framework for modelling belief change. They define belief change systems, or BCSs, which describe how the agent's beliefs about the world change. An axiomatic characterisation of BCSs is provided in a logical language $\mathcal{L}$ containing two modal operators. The unary operator B captures the agent's beliefs, and $B\alpha$ is read *the agent believes α* ; the binary operator $>$ represents the change of the agent's beliefs, and $\alpha > \beta$ is read *after learning α , the agent will be in an epistemic state satisfying β* . A special class of belief change systems, preferential belief change systems, is also investigated. In this case, an epistemic state can be identified with a set of interpretations of $\mathcal{L}$ together with a preference ordering the interpretations. An agent believes α in an epistemic state s exactly if α is true in all worlds considered possible at s and

the agent believes β after learning α in epistemic state s exactly if β is true in all the minimal⁸ interpretations that satisfy α .

One of the consequences of this approach lies in the fact that it investigates subtle differences between revision and update and better states when update is appropriate. In particular, concerning conditionals, since the truth value of conditionals depends on the agent's current belief state, it seems inappropriate to assume that it persists when that state changes. As mentioned in section 7, revision is supposed to deal with formulae that represent static worlds, while update is more appropriate for formulae whose truth depends on the current state.

12 Conclusion

It is worth noting that many approaches to revision were developed in the field of artificial intelligence in the last fifteen years, as illustrated in the present survey. Among the logical approaches, the AGM paradigm, in which revision is interpreted as beliefs change, has become a standard. In this approach an epistemic state is represented by a beliefs set, i.e. a deductively closed set of sentences of a logical language, which represents the agent's current beliefs. The KM reformulation of the AGM paradigm provides a representation theorem which specifies that a revision operation satisfying the AGM postulates is equivalent to a set of total pre-orders $\leq_\Psi$, where each pre-order corresponds to an epistemic state Ψ and is used for the revision of this state in the presence of new information. The principle of minimal change induces total pre-orders between formulae in the syntactic approaches, and total pre-orders between interpretations in the semantic approaches.

Coming back to the survey, in epistemic entrenchment relations, only the more entrenched formulae are kept. Nebel's approach uses stratified knowledge bases; the formulae with higher priority are preferred. In the syntactic possible worlds approach, the total pre-order stems from the closeness between worlds, defined by means of set inclusion. Groves' approach uses a system of spheres totally ordered by inclusion. In Dalal's revision the total pre-order between interpretations relies on the minimal Hamming distance between interpretations. Borgida constructs total pre-order between interpretations, based on a minimal set of propositional variables on which two interpretations differ. Ranking of the interpretations is used to define a total pre-order in Spohn's approach. Total pre-order is induced by the weighting of interpretations in Papini and Rauzy's approach. The representation theorem focuses on the minimum in the total pre-order and according to the interpretation of the principle of minimal change we get the same or a different result of revision, as illustrated in the reference example.

Although very elegant, this framework does not support iterated revision because the underlying preference relation between beliefs is lost in the process of change. In fact, in this approach an epistemic state is represented by a belief set; however, an epistemic state does not only consist of the agent's current beliefs but also encodes the strategy that the agent uses to modify his beliefs after learning a new piece of information. A revision operator satisfying the AGM postulates is equivalent to a set of pre-orders $\leq_\Psi$; however, the total pre-orders associated with two successive epistemic states are not related – the only requirement is that these pre-orders are faithful.

From a computational point of view, efficient and incremental methods with reasonable complexity have to be designed.⁹ In this case, one has to deal with finite knowledge bases. It is to be noted that no method is unanimously adopted in the community. Other authors try to define a more general formal framework in order to represent both revision and update as special cases of a more general change operation. Another research direction, in the field of revision, is that of iterated revision. In most realistic applications such as, for instance, machine learning, the revision process has not to be designed in a single step, but this is a recurrent process which evolves with time. Some authors, taking a theoretical point of view, try to define a suitable theoretical framework in order to

⁸According to the preference ordering interpretations.

⁹Complexity of revision operations is not dealt with in this survey. See Eiter & Gottlob (1992).

formulate postulates analogous to AGM postulates. Other authors, taking a computational point of view, focus on the implementation of incremental methods.

References

- Alchourron, C, Gärdenfors, P and Makinson, D, 1985, "On the Logic of Theory Change: Partial Meet Functions for Contraction and Revision" *Journal of Symbolic Logic* **50** 510–530.
- Berlandier, P and Neveu, B, 1994, "Maintaining arc consistency through constraint retraction" *Proceedings of the 6th I. E. E. Int. Conf. on Tools with Artificial Intelligence* SPAN.
- Bessière, C, 1991, "Arc-consistency in dynamic constraint satisfaction problem" *Proceedings of AAAI-91* 221–226.
- Borgida, A, 1985, "Language features for flexible handling of exceptions in information systems" *ACM Transactions on DataBase System* 563–603.
- Bouquet, F and Jegou, P, 1994, "Solving over-constrained CSP using weighted OBDD" *Proceedings of ECAI94* SPAN.
- Boutilier, C, 1993, "Revision sequences and nested conditionals", *Proceedings of IJCAI-93* 519–531.
- Boutilier, C, 1995, "Generalised update: belief change in dynamics settings" *Proceedings of IJCAI-95* 1550–1557.
- Bryant, R, 1992, "Symbolic boolean manipulation with ordered binary decision diagrams" *ACM Computing Surveys* 293–318.
- Dalal, M, 1988, "Investigations into theory of knowledge base revision" *Proceedings of AAAI-88* 475–164.
- Darwiche, A and Pearl, J, 1997, "On the logic of iterated revision" *Artificial Intelligence* **89** 1–29.
- De Kleer, J, 1986, "An assumption-based TMS" *Artificial Intelligence* **28** 127–162.
- Dechter, R and Dechter, A, 1988, "Belief maintenance in dynamic constraint networks" *Proceedings of AAAI-88* 37–42.
- Doyle, J, 1979, "A truth maintenance system" *Artificial Intelligence* **12** 231–272.
- Dubois, D, Lang, J and Prade, H, 1994, "Automated reasoning using possibilistic logic" *IEEE Transactions On Knowledge and Data Engineering* **6** 64–71.
- Eiter, T and Gottlob, G, 1992, "On the complexity of propositional knowledge base revision" *Artificial intelligence* **57** 227–270.
- Fagin, R, Kuper, GM, Ullman, JD, and Vardi, MY, 1986, "Updating logical databases" *Advances in Computing Research* **3** SPAN.
- Fagin, R, Ullman, JD, and Vardi, MY, 1983, "On the semantic of updates in databases" *Proceedings of the 2nd ACM Symposium on Principles of Data Base Systems* 352–365.
- Friedman, N and Halpern, J, 1994a, "Conditional logics of belief change" *Proceedings of the AAAI-94* 915–921.
- Friedman, N and Halpern, J, 1994b, "A knowledge base for belief change, part II: revision and update" *Proceedings of 4th Principles of Knowledge Representation and Reasoning Conference* 190–201.
- Friedman, N and Halpern, J, 1996, "Belief revision: a critique" *Proceedings of 5th Principle of Knowledge Representation and Reasoning Conference* 421–431.
- Fuhrmann, A, 1991, "Belief revision" *Proceedings of the 3rd European Summer School in Language, Logic and Information* 1–36.
- Gärdenfors, P, 1988, *Knowledge in Flux: Modeling the Dynamics of Epistemic States* MIT Press.
- Ginsberg, ML and Smith, DE, 1988, "Reasoning about action -I: a possible models approach; II : the qualification problem" *Artificial Intelligence* **35** 165–195 and 311–342.
- Girard, JY, 1987, "Linear logic" *Theoretical Computer Science* **50** 1–102.
- Grove, A, 1988, "Two modellings for theory change" *Journal of Philosophical Logic* **17** 157–170.
- Harper, W. L, 1975, "Rational belief change, Popper functions and counterfactuals" *Synthese* **30** SPAN.
- Jeannicot, S, Oxusoff, L and Rauzy, A, 1988, "Evaluation sémantique en calcul propositionnel" *Revue d'intelligence artificielle* **2** 41–60.
- Katsuno, H and Mendelzon, A, 1991a, "On the difference between updating a knowledge base and revising it" *Proceedings of the 2nd Principles of Knowledge Representation and Reasoning Conference* 387–394.
- Katsuno, H and Mendelzon, A, 1991b, "Propositional knowledge base revision and minimal change", *Artificial Intelligence*, **52** 263–294.
- Lehmann, D, 1995, "Belief revision revisited" *Proceedings of IJCAI-95* 1534–1539.
- Levi, I, 1980, *The Enterprise of Knowledge* MIT Press.
- Levy, F, 1994, "A survey of belief revision and updating in classical logic" *International Journal of Intelligent Systems* **9** 29–59.
- Lewis, D, 1973, *Counterfactuals* Harvard University Press.
- Lewis, DK, 1976, "Probabilities of conditionals and conditional probabilities" *Philosophical Review* **85** 297–315.

- Makinson, D, 1993, "Five faces of minimality" *Studia Logica* 52 339–379.
- Makinson, D and Gärdenfors, P, 1991, "The relations between the logic of theory change and nonmonotonic logic" *Logic of Theory Change: Lecture notes in Artificial Intelligence* 465 185–205.
- McCarthy, J, 1980a, "Application of circumscription to formalising common sense knowledge" *Artificial Intelligence* 28 89–116.
- McCarthy, J, 1980b, "Circumscription a form of non-monotonic reasoning" *Artificial Intelligence* 13 27–39.
- Montanari, U, 1974, Network of constraints: fundamental properties and applications to picture processing *Information Science* 7 95–132.
- Nebel, B, 1991, "Belief revision and default reasoning: syntax-based approach" *Proceedings of the 2nd Knowledge Representation and Reasoning Conference* 417–427.
- Papini, O, 1996a, "Calcul propositionnel étendu: un cadre pour la révision multiple" *Rapport interne* 1–34.
- Papini, O, 1996b, "Révision des connaissances" *Proceedings of JFPLC'96* 217–237.
- Papini, O and Rauzy, A, 1995, "Revision in extended propositional calculus" *Proceedings of ECSQARU'95* 328–335.
- Pearl, J, 1988, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference* Morgan Kaufmann.
- Reiter, R, "A logic for default reasoning" *Artificial Intelligence* 13 81–132.
- Schiex, T and Verfaillie, G, 1994, "Stubbornness : an enhancement scheme for backjumping and nogood recording" *Proceedings of the 12th ECAI* 165–169.
- Shoham, Y, 1988, *Reasoning about Change: Time and Causation from the Standpoint of Artificial Intelligence* MIT Press.
- Sombe, Lea, 1992, "Révision de bases de connaissances" *Actes des 4èmes journées nationales du PRC-GDR I A* 207–235.
- Sombe, Lea, 1994, *Revision and Updating In Knowledge Bases* J. Wiley.
- Spohn, W, 1987, "Ordinal conditional functions: a dynamic theory of epistemic states" *Causation in Decision, Belief Change, and Statistics* 2 105–134.
- "Sur Le Collectif du projet inter -PRC. Gestion de l'évolutif et de l'incertain dans une base de connaissances", 1995a, *Actes des 5ièmes journées nationales du PRC-GDR Intelligence Artificielle* 77–121.
- "Sur Le Collectif sur les CSP. Autour du problème de satisfaction de contraintes", 1995b, *Actes des 5ièmes journées nationales du PRC-GDR Intelligence Artificielle* 159–178.
- Vempaty, N, 1992, "Solving constraint satisfaction problems using finite state automata", *Proceedings of the AAAI-92* 453–458.
- von Wright, GH, 1965, "And Next" *Acta Philos. Fenn. Fasc.* 18 297–304.
- Willard, L and Li, YY, 1990, "The revised Gärdenfors postulates and update semantics" *Proceedings of the 3rd International Conference on Database Theory SPAN*.
- Williams, MA, 1994, "Transmutations of knowledge systems" *Proceedings of 4th Principles of Knowledge Representation and Reasoning Conference* 619–629.
- Williams, MA, 1995, "Iterated theory base change:a computational model" *Proceedings of IJCAI-95* 1541–1547.
- Winslett, M, 1988, "Reasoning about action using a possible models approach" *Proceedings of AAAI-88* 89–93.