

A general framework for revising belief bases using qualitative Jeffrey's rule

Salem Benferhat¹ and Didier Dubois² and Henri Prade² and Mary-Anne Williams³

¹ CRIL-CNRS, UMR 8188, Faculté Jean Perrin, Université d'Artois, Rue Jean Souvraz, 62307 Lens France

² IRIT - Université Paul Sabatier, 118 route de Narbonne 31062 Toulouse cedex 09 France

³ Innovation and Enterprise Research Laboratory University of Technology, Sydney NSW 2007 Australia

Abstract

Intelligent agents require methods to revise their epistemic state as they acquire new information. Jeffrey's rule, which extends conditioning to uncertain inputs, is used to revise probabilistic epistemic states when new information is uncertain. This paper analyses the expressive power of two possibilistic counterparts of Jeffrey's rule for modeling belief revision in intelligent agents. We show that this rule can be used to recover most of the existing approaches proposed in knowledge base revision, such as adjustment, natural belief revision, drastic belief revision, revision of an epistemic by another epistemic state. In addition, we also show that that some recent forms of revision, namely reinforcement operators, can also be recovered in our framework.

Introduction

An intelligent agent's information is often uncertain, inconsistent and incomplete. It is then crucially important to define mechanisms to manage it in response to focusing on a specific problem or in response to the acquisition of new, possibly conflicting, information. The term information covers a broad range of entities such as knowledge, perceptions, beliefs, expectations, preferences, or causal relations. It can describe agents' view of the world, itself, its actions and its understanding of changes.

During the past twenty years, many approaches have been proposed to address the problem of belief change from the axiomatic point of view (e.g., (Gärdenfors 1988), (Darwiche & Pearl 1997)), from the semantics point of view (e.g., (Williams 1994), (Boutilier 1993), (Thielscher 2005)) and from the computational point of view ((Nebel 1994), (Benferhat *et al.* 2002)).

Due to lack of space, this paper only focuses on the semantics of belief revision in the framework of possibility theory. The basic object in possibility theory is a possibility distribution, which is a mapping from the set of classical interpretations to an ordered structure, usually the interval $[0, 1]$. A possibility distribution rank-orders the potential states of the real world according to their level of plausibility, and represents the information available to an agent.

The revision of a possibility distribution can be viewed as a so-called "transmutation" (Makinson 1994) that modifies the ranking of interpretations so as to give priority to the input information. In particular, two forms of possibilistic

revision, called possibilistic revision with partial epistemic states, are investigated as the counterparts of Jeffrey's rule of revision in probability theory. These two forms of possibilistic revision consist in modifying a possibility distribution π with a set of weighted, mutually exclusive formulas $\mu = \{\phi_i, a_i\}$, called partial epistemic states, which express that the possibility of ϕ is equal to a_i . These two forms of revision come down to modifying the possibility π such that each formula ϕ_i is accepted with the prescribed degree a_i . The new degrees a_i 's may be either a constant determined for example by an expert, or a function defined for instance with respect to the original possibility degree associated with ϕ_i .

This paper first extends the natural properties described in (Benferhat *et al.* 2002) in order to take into account the new form of the input, namely a partial epistemic state. Then we present two definitions of possibilistic revision operators that naturally extend the two forms of conditioning that have been defined in the possibility theory framework. We also compare possibilistic revision with the counterpart of Jeffrey's rule of conditioning. In its second half, the paper shows that most of existing belief revision operators can be recovered by one of the two forms of possibilistic revision with respect to partial epistemic states.

But first in order to establish the new results, we need to restate the necessary background on possibility theory.

Possibilistic representations of epistemic states

Let L be a finite propositional language with formulas ϕ or ψ . \models denotes the (semantical) classical consequence relation. Ω is the set of classical interpretations, and $[\phi]$ is the set of classical models of ϕ .

An epistemic agent is a special kind of intelligent agent, one that at any given moment in time is in a specific epistemic state, e.g. it will have a set of current beliefs which are crafted from its background knowledge, conceptual understanding and its (internal and external) perceptions.

We take the traditional interpretation of beliefs and epistemic states and view an *epistemic state* as a set of beliefs where a belief is a relation between an epistemic agent and an object of belief represented as a logical sentence or a proposition.

There are several common representations of epistemic states such as : well ordered partitions of Ω , probabilistic

epistemic states, Grove's systems of spheres, Spohn's Ordinal Conditional Functions (OCF), etc. Throughout this paper we use a general representation of a total preorder, namely a possibility distribution π , which is a mapping from Ω to the interval $[0,1]$.

Indeed a possibility distribution can be used for representing any total preorder, and any operator on a total preorder on Ω can be translated into an operator on a possibility distribution to obtain the same outcome. We will identify operators that require the full power of the $[0,1]$ scale.

Given an interpretation $\omega \in \Omega$, $\pi(\omega)$ represents the degree of compatibility of ω with the available information (or beliefs) about the real world. $\pi(\omega) = 0$ means that the interpretation ω is impossible, and $\pi(\omega) = 1$ means that nothing prevents ω from being the real world. Interpretations ω where $\pi(\omega) = 1$ are considered to be expected (they are not at all surprising). When $\pi(\omega) > \pi(\omega')$, ω is a preferred candidate to ω' for being the real state of the world. The less $\pi(\omega)$ the less plausible ω or the more different it is to the current world. A possibility distribution π is said to be *normal* if $\exists \omega \in \Omega$, such that $\pi(\omega) = 1$, in other words if at least one interpretation is a fully plausible candidate for being the actual world.

Given a possibility distribution π , the possibility degree of formula ϕ is defined as:

$$\Pi_{\pi}(\phi) = \max\{\pi(\omega) : \omega \in [\phi]\}.$$

It evaluates the extent to which ϕ is consistent with the available information expressed by π . When there is no ambiguity, we simply write $\Pi(\phi)$ instead of $\Pi_{\pi}(\phi)$. Note that $\Pi(\phi)$ is evaluated under the assumption that the situation where ϕ is true is as normal as can be (since $\Pi(\phi)$ reflects the maximal plausibility of a model of ϕ).

Given a possibility distribution π , the semantic determination of the content of an epistemic state denoted by $\text{content}(\pi)$, is obtained by considering all sentences which are more plausible than their negation, namely:

$$\text{content}(\pi) = \{\phi : \Pi(\phi) > \Pi(\neg\phi)\}.$$

Namely, $\text{content}(\pi)$ is a classical theory whose models are the interpretations having the highest degrees in π . When π is normalized, models of $\text{content}(\pi)$ are interpretations which are completely possible, namely $[\text{content}(\pi)] = \{\omega : \pi(\omega) = 1\}$. The sentence ϕ belongs to $\text{content}(\pi)$ when ϕ holds in all the most normal or plausible situations (hence ϕ is expected, or accepted as being true).

Lastly, given a formula ϕ , two different types of conditioning (Dubois & Prade 1998) have been defined in possibility theory (when $\Pi(\phi) > 0$):

- In an ordinal setting, we have:

$$\begin{aligned} \pi(\omega \mid_m \phi) &= 1 \text{ if } \pi(\omega) = \Pi(\phi) \text{ and } \omega \models p \\ &= \pi(\omega) \text{ if } \pi(\omega) < \Pi(\phi) \text{ and } \omega \vdash p \\ &= 0 \text{ if } \omega \notin [p]. \end{aligned} \quad (1)$$

This is the definition of *minimum-based conditioning*.

- In a numerical setting, we get:

$$\begin{aligned} \pi(\omega \mid \cdot \phi) &= \frac{\pi(\omega)}{\Pi(\phi)} \text{ if } \omega \models p \\ &= 0 \text{ otherwise} \end{aligned} \quad (2)$$

This is the definition of *product-based conditioning*.

These two definitions of conditioning satisfy an equation of the form

$$\forall \omega, \pi(\omega) = \pi(\omega \mid \phi) \oplus \Pi(\phi)$$

which is similar to Bayesian conditioning, where \oplus is min and the product respectively. The rule based on the product is much closer to genuine Bayesian conditioning than the qualitative conditioning defined from the minimum which is purely based on the comparison of levels; product-based conditioning requires more of the structure of the unit interval. Besides, when $\Pi(\phi) = 0$, $\pi(\omega \mid_m \phi) = \pi(\omega \mid \cdot \phi) = 1, \forall \omega$, by convention.

Iterated semantic revision in possibilistic logic

Belief revision results from the effect of accepting a new piece of information called the input information. In this paper, it is assumed that the current epistemic state (represented by a possibility distribution), and the input information, do not play the same role. The input must be incorporated in the epistemic state. In other words, it takes priority over information in the epistemic state. This asymmetry is expressed by the way the belief change problem is stated, namely the new information alters the epistemic state and not conversely. This asymmetry will appear clearly at the level of belief change operations. This situation is different from the one of information fusion from several sources, where no epistemic state dominates *a priori*. In this context, the use of symmetrical rules is natural especially when the sources are equally reliable.

Jeffrey's rule for revising probability distributions

Reasoning in the presence of new observations is a fundamental issue in reasoning with uncertainty and imprecision. In probability theory, there is a natural method for achieving this task using Jeffrey's rule (Jeffrey 1965). This rule is proposed for revising probability distributions based on the probability kinematics principle whose objective is minimizing change. In this method, beliefs are represented as a probability distribution.

Jeffrey's rule (Jeffrey 1965) provides an effective means to revise a probability distribution p to p' given uncertainty bearing on a set of *mutually exclusive* and *exhaustive* events ϕ_i . Note that when speaking of events, ϕ is short for $[\phi]$. The uncertainty is given in the form of pairs (ϕ_i, a_i) with:

$$P'(\phi_i) = a_i. \quad (3)$$

Jeffrey's method relies on the fact that although there is uncertainty about events ϕ_i , conditional probability of any event $\psi \subseteq \Omega$ given any uncertain event ϕ_i remains the same in the original and the revised distributions. Namely,

$$\forall \phi_i, \forall \psi, P(\psi \mid \phi_i) = P'(\psi \mid \phi_i). \quad (4)$$

The underlying interpretation of revision implied by the constraint of Equation 4 is that the revised probability distribution p' must not change conditional probability degrees

of any event ϕ given uncertain events ϕ_i . In the probabilistic framework, applying Bayes rule then marginalization allows revision of the possibility degree of any event ψ in the following way:

$$P'(\psi) = \sum_{\phi_i} P'(\phi_i) * \frac{P(\psi, \phi_i)}{P(\phi_i)}. \quad (5)$$

The revised probability distribution p' (known as Jeffrey's rule of conditioning) is the unique distribution that satisfies (3) and (4) (see(Chan & Darwiche 2005)).

Two forms of possibilistic revision based on Jeffrey's rule

The possibilistic counterpart of Jeffrey's rule was introduced in (Dubois & Prade 1991) (see also (Dubois & Prade 1997)), without emphasizing the probability kinematics condition (4) however. There are two natural ways to define a possibilistic revision based on Jeffrey's rule, which naturally extend the two forms of conditioning that exist in possibility theory.

Note that most existing works on belief revision (both from semantics and axiomatics perspectives) assume that the input information is either a propositional formula, or an epistemic state (namely a possibility distribution).

Defining a possibilistic revision based on Jeffrey's rule allows us to define a general framework where the input is a compact partition of the set of interpretations. Namely, the input is of the form $\mu = \{(\phi_i, a_i) \mid i = 1, m\}$ where the ϕ_i 's are pairwise mutually exclusive formulas. The only requirement is that there exists at least one a_j such that $a_j = 1$. In the following, μ will be called a partial epistemic state. It is partial in the sense that letting $\Pi'_{(\pi')}(\phi_i) = a_i$ does not amount to the full specification of π' over the models of ϕ_i .

Let us first discuss some natural properties of the revision of a possibility distribution π and a new input information $\mu = \{(\phi_i, a_i) \mid i = 1, m\}$ to a new possibility distribution denoted by $\pi' = \pi(\cdot | \mu)$. Natural properties for π' are:

A₁ : π' should be normalized,

A₂ : $\forall(\phi_i, a_i) \in \mu, \Pi'(\phi_i) = a_i$.

A₃ : $\forall\omega, \omega' \in [\phi_i]$ then: $\pi(\omega) \geq \pi(\omega')$ then $\pi'(\omega) \geq \pi'(\omega')$,

A₄ : If for all $\phi_i, \Pi(\phi_i) = a_i$ then $\forall\omega \in [\phi_i] : \pi(\omega) = \pi'(\omega)$,

A₅ : If $\pi(\omega) = 0$ then $\pi'(\omega) = 0$.

A₁ means that the new epistemic state is consistent. **A₂** confirms that the input (ϕ, a) is interpreted as a constraint which forces π' to satisfy:

$$\Pi'(\phi_i) = a_i.$$

A₃ means that the new possibility distribution should preserve the previous relative order (in the wide sense) between models of each ϕ_i . A stronger version of **A₃** can be defined:

A₃' : $\forall\omega, \omega' \in [\phi_i]$ then: $\pi(\omega) > \pi(\omega')$ iff $\pi'(\omega) > \pi'(\omega')$,

A₃' clearly extends **CR₁**, **CR₂** proposed in (Darwiche & Pearl 1997). **A₄** means that when all new beliefs ϕ_i are accepted at their prescribed levels a_i then revision does not affect π . **A₅** stipulates that impossible worlds remain impossible after revision. Note that there are no further constraints which relate models of different ϕ_i in the new epistemic state.

The previous properties **A₁**–**A₅** do not guarantee a unique definition of conditioning.

A₃ suggests that the possibilistic revision process can be achieved using several parallel changes with a sure input: First, apply a conditioning (using equations 1 and 2) on each ϕ_i and in order to satisfy **A₂**, the distribution $\pi(\cdot | \neg\phi)$ is denormalized so as to satisfy $\Pi'(\phi_i) = a_i$. Therefore, revising with μ can be achieved using the following definition:

$$\forall\phi_i \in \mu, \forall\omega \models \phi_i, \pi(\omega | \mu) = a_i \oplus \pi(\omega | \oplus \phi_i) \quad (6)$$

where \oplus is either min or the product, depending on whether conditioning is based on the product or the minimum operator. When $\oplus =$ product (resp. min) the possibilistic revision will be simply called product-based (resp. minimum-based) conditioning with partial epistemic states.

The new degree of models of ϕ_i depends on the relative position of the *a priori* possibility degree of ϕ_i , and the prescribed posterior possibility degree of ϕ_i :

- If $\Pi(\phi_i) \geq a_i$ and when $\oplus =$ min, all interpretations that were originally more plausible than a_i , are forced to level a_i , which means that some strict ordering between models of ϕ_i may be lost. Hence **A₃'** is clearly not satisfied. When $\oplus =$ product, all plausibility levels are proportionally shifted down (to the level a_i).
- If $\Pi(\phi_i) < a_i$ the best models of ϕ_i are raised to level a_i . Moreover, when $\oplus =$ product, the plausibility levels of other models are proportionally shifted up (to level a_i).

Relationships with Jeffrey's kinematics properties

Another way to define possibilistic revision is to simply apply the counterpart of Jeffrey's rule of conditioning (Jeffrey 1965). Namely, given an initial possibility distribution π and a partial epistemic state $\mu = \{(\phi_i, a_i) \mid i = 1, m\}$ we need to find possibility distributions π' that satisfy:

$$\Pi'(\phi_i) = a_i. \quad (7)$$

and:

$$\forall\phi_i, \forall\psi, \Pi(\psi | \oplus \phi_i) = \Pi'(\psi | \oplus \phi_i), \quad (8)$$

where \oplus is either a minimum or a product. When \oplus is the product then we can show that the possibilistic revision given by (6) is the unique possibility distribution that satisfies (7) and (8). However, it is not the case when \oplus is the minimum.

Recovering existing belief revision frameworks

Standard possibilistic conditioning and adjustment

Clearly, possibilistic revision with partial epistemic states generalizes possibilistic conditioning with a propositional formula ϕ . Indeed, applying possibilistic revision given by (6) with a partial epistemic states $\mu = \{(\phi, 1), (\neg\phi, 0)\}$ gives exactly the same results if one applies equation (1) on ϕ when $\oplus = \min$ (resp. (2) for $\oplus = \text{product}$).

Similarly, possibilistic revision with uncertain input, which corresponds to adjustment (see (Benferhat *et al.* 2002)), is a particular case of possibilistic revision with a partial epistemic state, where the input is of the form $\mu = \{(\phi, 1), (\neg\phi, a)\}$.

Natural belief revision

Let $<_{initial}$ be a total pre-order on the set of epistemic states. Let ϕ be a new piece of information. We denote by $<_N$ the result of applying natural belief revision of $<_{initial}$ by ϕ . Natural belief revision of $<_{initial}$ by ϕ proposed in (Boutilier 1993), also hinted by Spohn (Spohn 1988), proceeds to minimal change of $<_{initial}$ by considering the most plausible models of ϕ in $<_{initial}$ to be the most plausible interpretation in $<_N$. More precisely, $<_N$ is defined as follows:

- $\forall \omega \in \min(\Omega, <_{initial}), \forall \omega' \in \min(\Omega, <_{initial}), \omega =_N \omega'$
- $\forall \omega \in \min(\Omega, <_{initial}), \forall \omega' \notin \min(\Omega, <_{initial}), \omega <_N \omega'$
- $\forall \omega \notin \min(\Omega, <_{initial}), \forall \omega' \notin \min(\Omega, <_{initial}), \omega <_N \omega'$ iff $\omega <_{initial} \omega'$.

To recover natural belief revision, first associate with $<_{initial}$ a compatible positive possibility distribution $\pi_{initial}$, defined by :

$\forall \omega, \omega' \in \Omega, \pi_{initial}(\omega) > \pi_{initial}(\omega')$ iff $\omega <_{initial} \omega'$. Such $\pi_{initial}$ always exists. Then let a be such that $1 > a > \max\{\pi(\omega) : \pi(\omega) \neq 1\}$. Then define $\pi_{<_N}(\cdot) = \pi_{input}(\cdot | \mu)$ where $\mu = \{(\phi, 1), (\neg\phi, a)\}$, $\pi_{input}(\cdot | \mu)$ is the result applying possibilistic revision given by equation (6) with $\oplus = \min$. Then we can show that $\pi_{<_N}$ indeed encodes natural belief revision, namely:

$$\forall \omega, \omega' \in \Omega, \pi_{<_N}(\omega) > \pi_{<_N}(\omega') \text{ iff } \omega <_N \omega'.$$

Drastic belief revision

Papini (Papini 2000), has considered a stronger constraint (also hinted by Spohn (Spohn 1988)) by imposing that each model of ϕ should be strictly preferred to each countermodel of $\neg\phi$, and moreover the relative ordering between models (resp. countermodels) of p should be preserved. More formally, let us denote by $<_D$ be result of applying drastic belief revision of $<_{initial}$ by ϕ . $<_D$ is defined as follows:

- $\forall \omega, \omega' \in [\phi], \omega <_D \omega'$ iff $\omega <_{initial} \omega'$.

¹a possibility distribution π is said to be positive if $\forall \omega, \pi(\omega) > 0$.

- $\forall \omega, \omega' \notin [\phi], \omega <_D \omega'$ iff $\omega <_{initial} \omega'$.
- $\forall \omega \in [\phi], \forall \omega' \notin [\phi], \omega <_D \omega'$.

To recover drastic belief revision, first associate with $<_{initial}$ a compatible positive possibility distribution π_{input} , as defined above. Let $\Delta(\phi) = \min\{\pi(\omega) : \omega \models p\}$, and a such that $a < \Delta(\phi)$.

Then define $\pi_{<_D}(\cdot) = \pi_{input}(\cdot | \mu)$ where $\mu = \{(\phi, 1), (\neg\phi, a)\}$, $\pi_{input}(\cdot | \mu)$ is the result applying possibilistic revision given by equation (6) with $\oplus = \text{product}$. Then we can show that $\pi_{<_D}$ indeed encodes drastic belief revision, namely:

$$\forall \omega, \omega' \in \Omega, \pi_{<_D}(\omega) > \pi_{<_D}(\omega') \text{ iff } \omega <_D \omega'.$$

A revision of epistemic state by epistemic state

In (Benferhat *et al.* 2000) (see also (Nayak 1994)) a revision of an epistemic state, denoted here by $<_{initial}$, by an input in the form of an epistemic state, denoted here by $<_{input}$, is defined. The obtained result is a new epistemic state, denoted by $<_L$ (L for lexicographic ordering), and defined as follows:

- $\forall \omega, \omega' \in \Omega$, if $\omega <_{input} \omega'$ then $\omega <_L \omega'$.
- $\forall \omega, \omega' \in \Omega$, if $\omega =_{input} \omega'$ then $\omega <_L \omega'$ iff $\omega <_{initial} \omega'$.

Namely, $<_L$ is obtained by refining $<_{input}$ by $<_{initial}$. For our purpose, we denote $\{E_0, \dots, E_n\}$ the partition of Ω induced by $<_{input}$. Namely:

- $\forall i, j \in \{0, \dots, n\}, E_i \cap E_j = \emptyset$, and $\bigcup_{i=0, \dots, n} E_i = \Omega$ (namely, E_i 's represent a partition of Ω)
- $\forall i \in \{0, \dots, n\}, \forall \omega, \omega' \in E_i, \omega =_{input} \omega'$,
- $\forall \omega, \omega' \in \Omega, \omega <_{input} \omega'$ iff $\omega \in E_i, \omega' \in E_j$ and $i < j$.

Let $\pi_{initial}$ and π_{input} be two positive possibility distributions associated respectively with $<_{initial}$ and $<_{input}$.

To recover this revision of an epistemic state by an epistemic state, first define $\pi_{<_L}(\cdot) = \pi_{input}(\cdot | \mu)$ where $\mu = \{(\phi_{E_i}, \epsilon^i) : i = 0, \dots, n\}$ is the result applying possibilistic revision given by equation (6) with $\oplus = \text{product}$. ϕ_{E_i} is a propositional formulas that exactly admits E_i as the set of its models. ϵ_i 's are infinitesimal (and by convention $\epsilon^0 = 1$).

Then we can show that $\pi_{<_L}$ indeed encodes $<_L$, namely:

$$\forall \omega, \omega' \in \Omega, \pi_{<_L}(\omega) > \pi_{<_L}(\omega') \text{ iff } \omega <_L \omega'.$$

Reinforcement operator

The last approach that we propose to recover is called a reinforcement operator, recently proposed in (Konieczny & Perez 2008). The idea is that a revision of $<_{initial}$ by a propositional formula ϕ only allows the improvement in plausibility of ϕ , namely the result makes ϕ "one unit" more plausible.

Let new epistemic state, denoted by $<_R$, obtained after reinforcing ϕ is defined as follows:

- The relative orderings between models (resp. countermodels) of ϕ is preserved.

- Let ω be a model of ϕ and ω' be a counter-model of ϕ :
 - if $\omega' = \omega$ then $\omega <_R \omega'$
 - if $\omega' <_{initial} \omega$ then if $\forall \omega'' \in [\phi]$ where $\omega' <_{initial} \omega''$ we have $\omega \leq \omega''$ then $\omega =_R \omega'$ otherwise $\omega' <_R \omega$
 - if $\omega <_{initial} \omega'$ then $\omega <_R \omega'$

To recover the reinforcement operator, we first define π_{input} to be a positive possibility distribution associated with $<_{initial}$, as defined above. Let $S = \{a_0 = 1, a_1, \dots, a_n\}$ be a finite scale $1 > a_1 > \dots > a_n > 0$ (n is at least equal to twice the number of different degrees in π_{input}). Define $pred(a_i) = a_{i-1}$ with by convention $pred(a_0) = 1$, and $succ(a_i) = a_{i+1}$ with by convention $Succ(a_n) = a_n$.

Let $a_i \in S$ be such that $a_i = \pi_{input}(\phi)$.

Define $\pi_{<_R}(\cdot) = \pi_{initial}(\cdot | \mu)(\cdot)$ where $\mu = \{(\phi, min(1, pred(\Pi(\phi))), (\neg\phi, succ(\Pi(\neg\phi)) \ominus \Pi(\phi))\}$ is the result of applying possibilistic revision given by equation (6) with $\oplus =$ product, and where “ $succ(\Pi(\neg\phi)) \ominus \Pi(\phi)$ ” is defined as equal to 1 if $\Pi(\neg\phi) > \Pi(\phi)$ and equal to $succ(\Pi(\neg\phi))$ otherwise.

Then we can show that $\pi_{<_R}$ indeed encodes $<_R$, namely:

$$\forall \omega, \omega' \in \Omega, \pi_{<_R}(\omega) > \pi_{<_R}(\omega') \text{ iff } \omega <_R \omega'.$$

Conclusion

The information held by an intelligent agent is typically uncertain, inconsistent and incomplete, consequently agents need sophisticated mechanisms to revise their epistemic states as they acquire new information over time because this new information may be in conflict with information in its epistemic state.

Due to the fundamental nature of the need to maintain an epistemic state that faithfully reflects an agents understanding there as been considerable scientific effort invested in developing effective belief revision mechanisms and strategies such as (Gärdenfors 1988), (Darwiche & Pearl 1997), (Williams 1994), (Nebel 1994), and (Benferhat *et al.* 2002).

In this paper we show how Jeffrey’s rule can be used to justify several key existing approaches to belief revision, then having established this sound relationship we show that reinforcement operators can be specified using our framework. Lastly, we propose a new form of belief revision where the input is only a partial representation of epistemic states using Jeffrey’s rule. All these methods can be used to enhance the belief management capabilities of intelligent agents.

References

Benferhat, S.; Konieczny, S.; Papini, O.; and Pérez, R. P. 2000. Iterated revision by epistemic states: axioms, semantics and syntax. In *Proc. of the 14th European Conf. on Artificial Intelligence (ECAI-00)*, 13–17. Berlin, Allemagne: IOS Press.

Benferhat, S.; Dubois, D.; Prade, H.; and Williams, M.-A. 2002. A practical approach to revising prioritized knowledge bases. *Studia Logica Journal* 70:105–130.

Boutilier, C. 1993. Revision Sequences and Nested Conditionals. In *Proc. of the 13th Inter. Joint Conf. on Artificial Intelligence (IJCAI’93)*, 519–531.

Chan, H., and Darwiche, A. 2005. On the revision of probabilistic beliefs using uncertain evidence. *Artificial Intelligence* 163:67–90.

Darwiche, A., and Pearl, J. 1997. On the logic of iterated revision. *Artificial Intelligence* 89:1–29.

Dubois, D., and Prade, H. 1991. *Updating with belief functions, ordinal conditional functions and possibility measures*. Uncertainty in Artificial Intelligence 6 (P. P. Bonissone, M. Henrion, L. N. Kanal, J. F. Lemmer, eds). Elsevier Science Publ. B.V. 311–329.

Dubois, D., and Prade, H. 1997. A synthetic view of belief revision with uncertain inputs in the framework of possibility theory. *Int. J. Approx. Reasoning* 17:295–324.

Dubois, D., and Prade, H. 1998. Possibility theory: qualitative and quantitative aspects. *Handbook of Defeasible Reasoning and Uncertainty Management Systems*. (D. Gabbay, Ph. Smets, eds.), Vol. 1: *Quantified Representation of Uncertainty and Imprecision* (Ph. Smets, ed.) 169–226.

Gärdenfors, P. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Bradford Books. Cambridge: MIT Press.

Jeffrey, R. C.. 1965. *The logic of decision*. Mc. Graw Hill, New York.

Konieczny, S., and Perez, R. P. 2008. Improvement operators. In *11th International Conference on Principles of Knowledge Representation and Reasoning (KR’08)*, 177–186.

Makinson, D. 1994. General patterns in nonmonotonic inference. In *Handbook of Logic in Artificial Intelligence and Logic Programming, Vol. 3* (D.M. Gabbay, C.J. Hogger, J.A. Robinson, D. Nute, eds.), 35–110. Oxford University Press,.

Nayak, A. 1994. Iterated belief change based on epistemic entrenchment. *Erkenntnis* 41:353–390.

Nebel, B. 1994. Base revision operations and schemes: semantics, representation, and complexity. In *Proceedings of the Eleventh European Conference on Artificial Intelligence (ECAI’94)*, 341–345.

Papini, O. 2000. Iterated revision operations stemming from the history of an agent’s observations. *Frontiers of Belief Revision* to appear.

Spohn, W. 1988. Ordinal conditiona functions: a dynamic theory of epistemic states. *Causation in Decision, Belief Change, and Statistics* 2:105–134.

Thielscher, M. 2005. Handling implicational and universal quantification constraints in flux. In van Beek., ed., *Proceedings of the International Conference on Principle and Practice of Constraint Programming (CP)*, volume 3709 of LNCS, 667–681. Sitges, Spain: Springer.

Williams, M. A. 1994. Transmutations of Knowledge Systems. In Doyle, J., and al. Eds., eds., *Inter. Conf. on principles of Knowledge Representation and reasoning (KR’94)*, 619–629. Morgan Kaufmann.