Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (2008)

A Semantic Approach for Iterated Revision in Possibilistic Logic

Guilin Qi
AIFB
Universität Karlsruhe
D-76128 Karlsruhe, Germany
gqi@aifb.uni-karlsruhe.de

Abstract

In this paper, we propose a new approach for iterated revision in possibilistic logic by applying a one-step revision operator. We first argue that the set of KM postulates for revision is too strong to define a practical one-step revision operator and some of them should be weakened. We then present a semantic approach for iterated revision in possibilistic logic using a one-step revision operator. The computation of the semantic approach is given. We show that our revision approach satisfies almost all the DP postulates for iterated revision and some other important logical properties.

Introduction

An intelligent agent often confronts the problem of revising her beliefs upon learning new information. This problem is called belief revision and has attracted much attention in the last decades. In their pioneer work, Alchourrón, Gärdenfors, and Makinson (AGM for short) proposed a set of postulates for characterizing a rational belief revision operator (see (Gärdenfors 1988)). In AGM's work, an agent's epistemic state is represented by a *belief set* and the revision process is controlled by a ranking on the belief set, called *epistemic entrenchment*. Later on, Katsuno and Mendelzon in (Katsuno & Mendelzon 1992) reformulated AGM postulates for revision in propositional logic and then showed that very few existing revision operators satisfy all the reformulated AGM postulates.

It has been argued that AGM postulates are too weak to characterize incremental adaption of beliefs and are applied to one-step revision operators only. Therefore, additional postulates for iterated belief revision have been proposed by Darwiche and Pearl (DP for short) in (Darwiche & Pearl 1997) and many others (such as (Nayak, Pagnucco & Peppas 2003) and (Jin & Thielscher 2007)). A concrete revision operator that satisfies all the AGM postulates and DP postulates was given in (Darwiche & Pearl 1997), where Spohn's *ordinal conditional functions* (OCF for short) are used to represent agent's epistemic states. This operator was revised in (Jin & Thielscher 2007) to satisfy an additional postulate. There are some other revision operators in OCF, such as Adjustment and Maxi-Adjustment (Williams 1994; Williams 1996).

It has been shown in (Benferhat et. al. 2002) that there is a correspondence relationship between OCF framework and possibilistic logic. Possibilistic logic is a weighted logic which is flexible for dealing with both inconsistency and uncertainty. In possibilistic logic, semantically, an epistemic state is represented by a *possibility distribution*, which is a mapping from the set of all interpretations to interval [0,1], and syntactically, it is represented by a *possibilistic knowledge base*, which is a finite set of weighted formulas. In (Benferhat et. al. 2002), the authors presented some syntactical methods for revising prioritized knowledge bases in the framework of possibilistic logic and OCF framework, that are semantically meaningful. The input for revision can be either a totally reliable input or an uncertain input of the form (p, a), where $a < 1$. We consider only a totally reliable input in this paper. In this case, the methods in (Benferhat et. al. 2002) do not satisfy DP postulates.

In this paper, we propose a new approach for iterated revision in possibilistic logic. We first argue that the set of postulates for characterizing a rational revision operator in propositional logic is too strong and some of them need to be weakened to characterize existing one-step revision operators. We then present a semantic approach for revising a possibilistic knowledge base by iteratively applying a one-step revision operator that results in a *normal* possibility distribution, and we provide its syntactical computation. The logical properties of our approach are considered. We finally present related work and conclusion.

Preliminaries

We consider a propositional language $\mathcal{L}_{PS}$ defined from a finite set of propositional variables PS and the usual connectives. The classical consequence relation is denoted as $\vdash$. An interpretation is a total function from PS to $\{0, 1\}$, denoted by a bit vector whenever a strict total order on PS is specified. Ω denotes the set of all interpretations. $Mod(\phi)$ denotes the set of models of a formula ϕ. Let M be a set of interpretations, $form(M)$ denotes the logical formula (unique up to logical equivalence) whose models are M. A *knowledge base K* is a finite set of propositional formulas which is *sometimes* identified with the conjunction of its elements. K is consistent if and only if $Mod(K) \neq \varnothing$. Two knowledge bases K_1 and K_2 are equivalent, denoted $K_1 \equiv K_2$, if and only if $Mod(K_1) = Mod(K_2)$.

Possibilistic logic

We briefly introduce possibilistic logic (see (Dubois, Lang, & Prade 1994) for more details). The semantics of possibilistic logic is based on the notion of a *possibility distribution* $\pi : \Omega \rightarrow [0,1]$. The possibility degree $\pi(\omega)$ represents the degree of compatibility (resp. satisfaction) of the interpretation ω with the available beliefs about the real world. A possibility distribution is said to be *normal* if $\exists \omega_0 \in \Omega$, such that $\pi(\omega_0) = 1$. From a *possibility distribution* π, two measures can be determined: the possibility degree of formula ϕ, $\Pi_\pi(\phi) = max\{\pi(\omega) : \omega \in \Omega, \omega \models \phi\}$ and the necessity degree of formula ϕ, $N_\pi(\phi) = 1 - \Pi_\pi(\neg\phi)$.

At the syntactic level, a *possibilistic formula*, is represented by a pair $(\phi,\ a)$, where ϕ is a propositional formula and a is an element of the semi-open real interval $(0,1]$, which means that the necessity degree of ϕ is at least equal to a, i.e. $N(\phi) \geq a$. Then uncertain or prioritized pieces of information can be represented by a *possibilistic knowledge base* which is a finite set of possibilistic formulas of the form $B = \{(\phi_i,\ a_i) : i = 1,...,n\}$. The classical base associated with B, denoted B^*, is defined as $B^* = \{\phi_i | (\phi_i,\ a_i) \in B\}$. A possibilistic knowledge base B is consistent if and only if its classical base B^* is consistent. Given a possibilistic knowledge base B, a unique *possibility distribution*, denoted π_B, can be obtained by the principle of minimum specificity (Dubois, Lang, & Prade 1994). For all $\omega \in \Omega$,

$$\pi_B(\omega) = \begin{cases} 1 & \text{if } \forall(\phi_i,\ a_i) \in B, \omega \models \phi_i, \\ 1 - max\{a_i | \omega \not\models \phi_i, (\phi_i, a_i) \in B\} & \text{otherwise.} \end{cases} \tag{1}$$

The a-cut (resp. strict a-cut) of a possibilistic knowledge base B is $B_{\geq a} = \{\phi_i \in B^* | (\phi_i, b_i) \in B \ and\ b_i \geq a\}$ (resp. $B_{>a} = \{\phi_i \in B^* | (\phi_i, b_i) \in B \ \ and\ b_i > a\}$). The *inconsistency degree* of B is: $Inc(B) = max\{a_i : B_{\geq a_i}\ is\ inconsistent\}$. That is, the inconsistency degree of B is the largest weight a_i such that the a_i-cut of B is inconsistent. Two possibilistic knowledge bases B and B' are said to be equivalent, denoted $B \equiv_s B'$, iff $\forall\ a \in (0,1],\ B_{\geq a} \equiv B'_{\geq a}$.

Definition 1 Let B be a possibilistic knowledge base. A formula ϕ is said to be a consequence of B to a degree a, denoted $B \vdash_\pi (\phi,\ a)$, iff (i)$B_{\geq a}$ is consistent; (ii) $B_{\geq a} \vdash \phi$; (iii)$\forall b > a$, $B_{\geq b} \not\vdash \phi$. ϕ is said to be a plausible consequence of B, denoted $B \vdash_P \phi$, iff there exists an $a > Inc(B)$, $B \vdash_\pi (\phi, a)$.

KM postulates for belief revision

Alchourrón, Gärdenfors and Markinson proposed a set of postulates for characterizing a rational one-step revision operator (see (Gärdenfors 1988)). In (Katsuno & Mendelzon 1992), AGM postulates for revision are rephrased in propositional logic as follows, where $\circ$ is a revision operator which is a function from a knowledge base[1] K and a formula μ to a new knowledge base denoted by $K \circ \mu$.

(R1) $K \circ \mu \vdash \mu$
(R2) If $K \wedge \mu$ is satisfiable then $K \circ \mu \equiv K \wedge \mu$
(R3) If μ is satisfiable then $K \circ \mu$ is also satisfiable
(R4) If $K_1 \equiv K_2$ and $\mu_1 \equiv \mu_2$ then $K_1 \circ \mu_1 \equiv K_2 \circ \mu_2$
(R5) $(K \circ \mu) \wedge \phi$ implies $K \circ (\mu \wedge \phi)$
(R6) If $(K \circ \mu) \wedge \phi$ is satisfiable then $K \circ (\mu \wedge \phi)$ implies $(K \circ \mu) \wedge \phi$

We call the rephrased AGM postulates as KM postulates for revision (or KM postulates). Explanations of these postulates can be found in (Katsuno & Mendelzon 1992).

A representation theorem has been given to establish the correspondence between the KM postulates and a *faithful assignment*. Let K be a knowledge base. A function that assigns to K a pre-order[2] over Ω, denoted $\preceq_K$, is said to be faithful if the following three conditions hold:

(1) If $\omega, \omega' \in Mod(K)$, then $\omega \prec_K \omega'$ does not hold.

(2) If $\omega \in Mod(K)$ and $\omega' \notin Mod(K)$, then $\omega \prec_K \omega'$ holds.

(3) If $K \equiv K'$, then $\preceq_K = \preceq_{K'}$.

Theorem 1 (Representation Theorem in (Katsuno & Mendelzon 1992)) A revision operator $\circ$ satisfies the postulates (R1)-(R6) if and only if there exists a faithful assignment that maps each knowledge base K to a total pre-order[3] $\preceq_K$ such that $Mod(K \circ \mu) = min(Mod(\mu), \preceq_K)$.

Revisit KM Postulates for Revision

In this section, we give a short discussion on the KM postulates. The purpose of this section is not to propose some new postulates for replacing KM postulates but to select some KM postulates or their weakened versions for characterizing some frequently used one-step revision operators. A typical revision operator satisfying all the KM postulates is the Dalal's revision operator $\circ_D$ which is defined as follows:

The distance between a knowledge base K (here, K is identified with the conjunction of its elements) and an interpretation ω is defined as $dist(K,\ \omega) = min_{\omega_i \in Mod(K)} dist(\omega, \omega_i)$, where $dist(\omega, \omega_i)$ is the Hamming distance between ω and ω_i. A total pre-order $\preceq_K$ is defined as $\omega_1 \preceq_K \omega_2$ if and only if $dist(Mod(K), \omega_1) \leq dist(Mod(K), \omega_2)$. Finally, Dalal's revision operator $\circ_D$ can be defined as: $Mod(K \circ_D \mu) = min(Mod(\mu), \preceq_K)$.

Another revision operator satisfying all the KM postulates, called amnesic revision operator (Rott 2000), denoted $\circ_a$, is defined as follows: $K \circ_a \mu = K \cup \{\mu\}$ if $K \cup \{\mu\}$ is consistent and μ otherwise. We will show later that even this revision operator will result in an interesting revision operator in possibilistic logic.

However, as shown in (Katsuno & Mendelzon 1992), many other revision operators do not satisfy all the KM postulates. It has been argued in many papers, such as (Nebel 1991; Nebel 1998; Hansson 1999; Jin & Thielscher 2007), that postulate (R4) which states the *principle of irrelevance of syntax* is not desirable from a pragmatic view. In this case,

[1] In (Katsuno & Mendelzon 1992), a knowledge base is identified with the conjunction of its elements.

[2] Recall that a pre-order $\preceq$ over Ω is a transitive and reflexive relation. As usual, $\omega \prec \omega'$ iff $\omega \preceq \omega'$ but $\omega' \not\preceq \omega$.

[3] A pre-order $\preceq$ over Ω is total if for every $\omega, \omega' \in \Omega$, either $\omega \preceq \omega'$ or $\omega' \preceq \omega$

a knowledge base should have different meaning from the conjunction of its elements and both should be treated differently. A weakening of (R4) is given in (Jin & Thielscher 2007): **(R4′)** If $\mu_1 \equiv \mu_2$ then $K \circ \mu_1 \equiv K \circ \mu_2$.

We give two typical examples of syntax-based revision operators (or formula-based revision operators). The first approach is called full meet revision operator and has been defined in several papers (Fagin, Ullman, & Vardi 1983; Ginsberg 1986; Nebel 1998) but in different forms. Let $(K \bot \mu)$ be the maximal subsets of K that are consistent with μ, i.e.,

$$K \bot \mu = \{A \subseteq K | A \not\vdash \neg\mu, \forall B \subseteq K,\ if\ A \subset B,\ then\ B \vdash \neg\mu\}.$$

Then, the full meet revision operator, denoted $\circ_F$, is defined as follows: $K \circ_F \mu = \bigvee (K \bot \mu) \cup \{\mu\}$.

Next, we introduce cardinality-maximizing revision operator, which has been defined in (Ginsberg 1986; Nebel 1998). Let $(K \bot_C \mu)$ be the cardinality-maximizing subsets of K that are consistent with μ, i.e.,

$$K \bot_C \mu = \{A \subseteq K | A \not\vdash \neg\mu, \forall B \subseteq K,\ |A| < |B|,\ then\ B \vdash \neg\mu\},$$

where $|A|$ denotes the cardinality of the set A. The cardinality-maximizing revision operator, denoted $\circ_{CM}$, is defined as follows: $K \circ_{CM} \mu = \bigvee (K \bot_C \mu) \cup \{\mu\}$.

In terms of postulates (R1)-(R6), the following holds[4].

Proposition 1 The full meet revision operator $\circ_F$ satisfies (R1), (R2), (R3), (R4′) and (R5). However, it does not satisfy (R4) and (R6) in general. The cardinality-maximizing revision operator $\circ_{CM}$ satisfies (R1), (R2), (R3), (R4′), (R5) and (R6). It does not satisfy (R4) in general.

According to Proposition 1, both operator $\circ_F$ and operator $\circ_{CM}$ satisfy (R1), (R2), (R3), (R4′) and (R5). Furthermore, operator $\circ_{CM}$ satisfies (R6). However, both operators do not satisfy (R4). It has been shown in (Katsuno & Mendelzon 1992) that (R6) is a condition which guarantees the pre-order $\preceq_K$ in Theorem 1 is a total pre-order and that this condition is too strong in some case and can be weakened to the following two conditions:

(R7) If $\psi \circ \mu_1 \vdash \mu_2$ and $\psi \circ \mu_2 \vdash \mu_1$ then $\psi \circ \mu_1 \equiv \psi \circ \mu_2$.
(R8) $(\psi \circ \mu_1) \wedge (\psi \circ \mu_2) \vdash \psi \circ (\mu_1 \vee \mu_2)$.

It was shown that the revision operators based on partial pre-orders are completely characterized by Conditions (R1)-(R5), (R7) and (R8). We have the following proposition.

Proposition 2 $\circ_F$ satisfies (R7) and (R8).

According to Propositions 1 and 2, both full meet revision operator and cardinality-maximizing revision operator can be equivalently defined by a pre-order over interpretations. Next, we define such pre-orders. Let K be a knowledge base and μ be a formula. Let $\omega^K = \{\phi \in K : \omega \models \phi\}$, then we can define two pre-orders $\preceq_F$ and $\preceq_{CM}$ as follows:

- $\omega_1 \preceq_F \omega_2$ iff $\omega_2^K \subseteq \omega_1^K$.
- $\omega_1 \preceq_{CM} \omega_2$ iff $|\omega_2^K| \leq |\omega_1^K|$.

[4] In (Nebel 1998), Nebel has checked logical properties of operators $\circ_F$ and $\circ_{CM}$ against AGM's original postulates defined on belief set. However, there is no work on checking which KM postulates they satisfy.

The following proposition shows that the operators $\circ_F$ and $\circ_{CM}$ can be defined by $\preceq_F$ and $\preceq_{CM}$ respectively.

Proposition 3 Let K be a knowledge base and μ be a formula. Then we have $K \circ_F \mu = min(Mod(\mu), \preceq_F)$ and $K \circ_{CM} \mu = min(Mod(\mu), \preceq_{CM})$.

As we have shown, three commonly used revision operators all satisfy Postulates (R1), (R2), (R3), (R4′), (R5), (R7) and (R8). This amounts to weaken (R4) and (R6). We argue that these postulates are more appropriate for characterizing a practical revision operator and call an operator satisfying them as a weak-AGM revision operator.

We give the following representation theorem.

Theorem 2 Revision operator $\circ$ is a weak-AGM revision operator iff there exists an assignment that maps each knowledge base K to a partial pre-order $\preceq_K$, which satisfies Conditions (1) and (2) for a faithful assignment, such that $Mod(K \circ \mu) = min(Mod(\mu), \preceq_K)$.

A Semantic Approach for Iterated Revision in Possibilistic Logic

An algorithm for semantic revision

Iterated revision in possibilistic logic has been discussed in (Benferhat et. al. 2002) and some revision operators were given. When the input formula μ is totally reliable and we want to revise B, the result of revision by the approaches in (Benferhat et. al. 2002) does not include the formulas whose weights are less than the inconsistency degree of $B \cup \{(\mu, 1)\}$. So too much information may be lost after revision if the inconsistency degree is high. This problem is called "drowning effect". To resolve the drowning effect, we use a weak-AGM revision operator to resolve inconsistency iteratively. Let $B = \{(\phi_1, a_1), ..., (\phi_n, a_n)\}$ be a possibilistic knowledge base. Let us rearrange the weights of formulas in B such that $b_1 > b_2 > ... > b_l > 0$, where b_i $(i = 1, ..., l)$ are all the distinct weights appearing in B. Let $\epsilon \in (0, 1)$ be a very small number such that $\epsilon < b_2$ (for example 0.0001).

In Algorithm 1, we first assign degree 0 to all the interpretations which do not satisfy the new information μ and then we go to the "while" loop. For a "while" loop such that $i \leq l$, we first collect all the formulas whose certainty degrees are equal to b_i and revise these formulas (i.e., S_i) by the formula whose models are M_i (i.e. $form(M_i)$) using operator $\circ$. The result of revision M_{i+1} is used for another "while" loop. For $\omega \in M_i \backslash M_{i+1}$ (i.e., ω does not satisfy the revised base), if $b_i = 1$, then $\pi(\omega) = 1 - b_i + \epsilon^{b_{i+1}+1}$; otherwise, $\pi(\omega) = 1 - b_i$. Note that when $b_i = 1$, the possibility degree of ω is increased by $\epsilon^{b_{i+1}+1}$. By doing this we can ensure that new evidence is absolutely more reliable than any information in the original knowledge base (see postulate (Rec) in Section *Logical Properties*).

According to Theorem 2, in each "while" loop, the result of revision is $M_{i+1} = min(M_i, \preceq_{S_i})$, where $\preceq_{S_i}$ is the preorder induced by S_i and the revision operator. Therefore, in each "while" loop, we change the information in each stratum S_i minimally to accommodate $form(M_i)$.

Let us define a lexicographic relation $\preceq_{lex,B,k}$ $(1 \leq k \leq n)$ as: $\omega \preceq_{lex,B,k} \omega'$ if and only if $\omega \preceq_{S_i} \omega'$ for all $i \leq k$ or $\exists i \leq k$

Algorithm 1: Algorithm for semantic revision

Data: a possibilistic knowledge base $B = \{(\phi_1, a_1), ..., (\phi_n, a_n)\}$; a formula μ; a weak-AGM revision operator $\circ$
Result: revised possibility distribution $\pi_{B,\mu,\circ}$
begin
  $M_1 := Mod(\mu)$
  i:=1
  for $\omega \not\models \mu$ **do**
    $\pi_{B,\mu,\circ}(\omega) := 0$
  while $i \leq l$ **do**
    $S_i := \{\phi_j : (\phi_j, a_j) \in B, a_j = b_i\}$
    $M_{i+1} := Mod(S_i \circ form(M_i))$
    for $\omega \in M_i \setminus M_{i+1}$ **do**
      if $b_i = 1$ **then**
        $\pi_{B,\mu,\circ}(\omega) := \epsilon^{b_{i+1}+1}$
      else
        $\pi_{B,\mu,\circ}(\omega) := 1 - b_i$
    $i := i + 1$
  for $\omega \in M_{l+1}$ **do**
    $\pi_{B,\mu,\circ}(\omega) := 1$
  return $\pi_{B,\mu,\circ}$
end

such that $\omega \prec_{S_i} \omega'$ and $\omega =_{S_j} \omega'$ for all $j < i$, where $\omega =_{S_j} \omega'$ iff $\omega \preceq_{S_j} \omega'$ and $\omega' \preceq_{S_j} \omega$. We give a representation theorem for the semantic revision approach:

Theorem 3 Given a possibilistic knowledge base B and a formula μ. Let b_i ($i = 1, ..., l$) be all the distinct weights appearing in B such that $b_i > b_{i+1}$ for every i. Suppose $\pi_{B,\mu,\circ}$ is returned by Algorithm 1. We then have (1) $\pi_{B,\mu,\circ}(\omega) = 1 - b_i$ iff $\omega \in min(Mod(\mu), \preceq_{lex,B,i-1})$ and $\omega \notin min(Mod(\mu), \preceq_{lex,B,i})$ for $1 < i \leq l$; and (2) $\pi_{B,\mu,\circ}(\omega) = 1$ iff $\omega \in min(Mod(\mu), \preceq_{lex,B,l})$.

Theorem 3 shows that our revision approach can be characterized by the lexicographic relation $\leq_{lex,B,k}$.

Example 1 Let $B = \{(p \vee r, 0.9), (p \vee \neg q, 0.8), (\neg r, 0.8), (q, 0.8)\}$ and $\mu = p \wedge r$. So there two distinct weights $b_1 = 0.9$ and $b_2 = 0.8$. There are eight interpretations: $\omega_1 = 000$, $\omega_2 = 001$, $\omega_3 = 010$, $\omega_4 = 011$, $\omega_5 = 100$, $\omega_6 = 101$, $\omega_7 = 110$, $\omega_8 = 111$. For example, ω_2 assigns 0 to both p and q and 1 to r. We now apply Algorithm 1 to revise B by μ. Suppose the full meet revision operator $\circ_F$ is used in the algorithm. It is clear that $\omega_i \not\models \mu$, for $i = 1, 2, 3, 4, 5, 7$, so $M_1 = \{\omega_6, \omega_8\}$ and $\pi(\omega_i) = 0$ for $i = 1, 2, 3, 4, 5, 7$. Let $i = 1$. $S_1 = \{p \vee r\}$. Since $(S_1) \cup \{\mu\}$ is consistent, we have $M_2 = Mod(S_1 \circ_F \mu) = \{\omega_6, \omega_8\}$. So $form(M_2) \equiv \mu$. Let $i = 2$. $S_2 = \{p \vee \neg q, \neg r, q\}$ and $S_2 \cup \{\mu\}$ is inconsistent. Since $\omega_6^{S_2} = \{p \vee \neg q\}$ and $\omega_8^{S_2} = \{p \vee \neg q, q\}$, we have $\omega_8 \prec_F \omega_6$. So $Mod(S_2 \circ_F \mu) = \{\omega_8\}$ and $\pi(\omega_6) = 0.2$. The "while" loop finishes. Finally, we have $\pi(\omega_8) = 1$.

Syntactical counterpart of semantic revision

In this section, we propose an algorithm to compute a possibilistic knowledge base corresponding to possibility distribution resulted from Algorithm 1.

Algorithm 2: Algorithm for syntactical revision

Data: a possibilistic knowledge base $B = \{(\phi_1, a_1), ..., (\phi_n, a_n)\}$; a formula μ; a weak-AGM revision operator $\circ$
Result: a possibilistic knowledge base $B_{\mu,\circ}$
begin
  $K_1 := \{\mu\}$
  i:=1
  $B_{\mu,\circ} := \{(\phi, 1) : \phi \in K_1\}$
  while $i \leq l$ **do**
    $S_i := \{\phi_j : (\phi_j, a_j) \in B, a_j = b_i\}$
    $K_{i+1} := S_i \circ \wedge_{\phi \in K_i} \phi$
    if $b_i = 1$ **then**
      $B_{\mu,\circ} := B_{\mu,\circ} \cup \{(\phi, 1 - \epsilon^{b_{i+1}+1}) : \phi \in K_{i+1} \setminus K_i\}$
    else
      $B_{\mu,\circ} := B_{\mu,\circ} \cup \{(\phi, b_i) : \phi \in K_{i+1} \setminus K_i\}$
    $i := i + 1$
  $B_{\mu,\circ} := B_{\mu,\circ} \cup \{(\phi, b_l) : \phi \in K_{l+1} \setminus K_l\}$
  return $B_{\mu,\circ}$
end

In the ith "while" loop of Algorithm 2, the set of formulas in B whose weights are b_i is revised by the conjunction of knowledge base obtained by the i-1th "while" loop using a weak-AGM revision operator. If $b_i \neq 1$, we assign weight b_i to formulas in the resulting knowledge base K_{i+1} and not in K_i. Otherwise, the weight of formulas in K_{i+1} and not in K_i is decreased to $1 - \epsilon^{b_{i+1}+1}$. In Example 1, the resulting knowledge base of revision is $B_{\mu,\circ} = \{(p \wedge r, 1), (p \vee r, 0.9), (p \vee \neg q, 0.8), (q, 0.8)\}$. Note that if we apply the approach in (Benferhat et. al. 2002), then the formulas $(p \vee \neg q, 0.8)$ and $(q, 0.8)$ are lost.

We give the correspondence between two algorithms.

Proposition 4 Given a possibilistic knowledge base B and a formula μ, suppose $\pi_{B,\mu,\circ}$ is the possibility distribution obtained by Algorithm 1 and $B_{\mu,\circ}$ is the possibilistic knowledge base obtained by Algorithm 2, then $\pi_{B_{\mu,\circ}}(\omega) = \pi_{B,\mu,\circ}(\omega)$ for all $\omega \in \Omega$.

It is clear that Algorithm 2 needs to apply the revision operator $\circ$ at most l times, so the complexity of our algorithm is not much harder than that of the revision operator. For example, suppose $\circ = \circ_D$, then it needs at most $[\mathcal{O}(log\ n)]$ calls to a NP oracle to generate a revised base.

Next, we show that some prioritized base revision operators can be computed by our algorithm. We follow the notations in (Delgrande, Dubois, & Lang 2006): Suppose $\succ$ is a strict order (i.e. a transitive and asymmetric binary relation) on a set X, then for any subset Y of X, we define the set of undominated elements of Y w.r.t. $\succ$, denoted $Max(\succ, Y)$, as: $Max(\succ, Y) = \{y \in Y | \text{there is no } z \in Y \text{ such that } z \succ y\}$.

Given a possibilistic knowledge base $B = \{(\phi_1, a_1), ..., (\phi_n, a_n)\}$, suppose b_i ($i = 1, ..., l$) are all the distinct weights of formulas appearing in B such that $b_1 > b_2 > ... > b_l > 0$. Then B can be represented as a stratified knowledge base $\Sigma_B = (S_1, ..., S_l)$, where $S_i = \{\phi : (\phi, a) \in B, a = b_i\}$. A subset Σ' of a stratified knowledge base Σ_B is also a stratified knowledge base. We denote $Cons(\Sigma_B, \mu)$ the set of all subsets of Σ_B which are consistent with μ, that is, the set of all stratified knowledge bases $\Sigma' = (S'_1, ..., S'_l)$ such that $\Sigma' \subseteq \Sigma_B$ and $\bigwedge(\Sigma') \wedge \mu$ is consistent, where $\bigwedge(\Sigma') = \wedge_{\phi \in \Sigma'} \phi$. We then have the following prioritized base revision operators:

discrimin (Brewka 1989; Nebel 1991) For $\Sigma' = (S'_1, ..., S'_l)$, $\Sigma'' = (S''_1, ..., S''_l) \in Cons(\Sigma_B, \mu)$, define $\Sigma'' \succ_{discrimin} \Sigma'$ iff $\exists i$ such that

(a) $S_i \cap S''_i \supset S_i \cap S'_i$, and

(b) for all $j < i$, $S_j \cap S''_j = S_j \cap S'_j$.

Then $\Sigma_B \circ_{discrimin} \mu =$
$\bigvee\{\bigwedge(\Sigma), \Sigma \in Max(\succ_{discrimin}, Cons(\Sigma_B, \mu))\} \wedge \mu$.

leximin (Benferhat et. al. 1993) For $\Sigma' = (S'_1, ..., S'_l)$, $\Sigma'' = (S''_1, ..., S''_l) \in Cons(\Sigma_B, \mu)$, define $\Sigma'' \succ_{leximin} \Sigma'$ iff $\exists i$ such that

(a) $|S_i \cap S''_i| > |S_i \cap S'_i|$, and

(b) for all $j < i$, $|S_j \cap S''_j| = |S_j \cap S'_j|$.

Then $\Sigma_B \circ_{leximin} \mu =$
$\bigvee\{\bigwedge(\Sigma), \Sigma \in Max(\succ_{leximin}, Cons(\Sigma_B, \mu))\} \wedge \mu$.

linear base revision (Nebel 1998) Let $\Sigma_B = (S_1, ..., S_l)$ and μ be a formula. The linear base revision, denoted $\circ_{linear}$, is defined inductively as follows: $\Sigma_B \circ_{linear} \mu = \{\mu\} \cup S'_1 \cup ... \cup S'_l$, where S'_i is defined by $S'_0 = \{\mu\}$, $S'_i = S_i$ if $S_i \cup \{\mu\} \cup \bigcup_{j=1}^{i-1} S'_j$ is consistent, $\varnothing$ otherwise, for $i \geq 1$.

Proposition 5 Let B be a possibilistic knowledge base and Σ_B its associated stratified knowledge base. Let $\circ_F$, $\circ_{CM}$ and $\circ_a$ be the full meet revision operator, the cardinality-maximizing revision operator and the amnestic revision operator respectively. We then have the following equations:

- $(B_{\mu,\circ_F})^* \equiv \Sigma_B \circ_{discrimin} \mu$,
- $(B_{\mu,\circ_{CM}})^* \equiv \Sigma_B \circ_{leximin} \mu$, and
- $(B_{\mu,\circ_a})^* \equiv \Sigma_B \circ_{linear} \mu$.

Logical Properties

In this section, we discuss the logical properties satisfied by our revision approach. We first consider generalized KM postulates for revision. Given possibilistic knowledge base B and weak-AGM revision operator $\circ$, we define $B \circ \mu = B_{\mu,\circ}$, that is, $B \circ \mu$ is the revised result of B by μ using Algorithm 2 and $\circ$.

Proposition 6 Let $\circ$ be a weak-AGM revision operator and B a possibilistic knowledge base. Then we have the following results:

(R^*_1) $B \circ \mu \vdash_\pi (\mu, 1)$.

(R^*_2) If $B^* \wedge \mu$ is consistent, then $(B \circ \mu)^* \equiv B^* \wedge \mu$.

(R^*_3) If μ is consistent, then $(B \circ \mu)^*$ is consistent.

(R^*_4) $(B \circ \mu)^* \wedge \phi \vdash (B \circ (\mu \wedge \phi))^*$.

(R^*_5) If $(B \circ \mu_1)^* \vdash \mu_2$ and $(B \circ \mu_2)^* \vdash \mu_1$, then $(B \circ \mu_1)^* \equiv (B \circ \mu_2)^*$.

(R^*_6) $(B \circ \mu_1)^* \wedge (B \circ \mu_2)^* \vdash (B \circ (\mu_1 \vee \mu_2))^*$.

Furthermore, if $\circ$ satisfies all the KM postulates, then it satisfies (R^*_1)-(R^*_6) when it is applied to revise a possibilistic knowledge base and the following two postulates:

(R^*_7) If $B_1 \equiv_s B_2$ and $\mu_1 \equiv \mu_2$, then $B_1 \circ \mu_1 \equiv_s B_2 \circ \mu_2$.

(R^*_8) If $(B \circ \mu)^* \wedge \phi$ is consistent, then $(B \circ (\mu \wedge \phi))^* \vdash (B \circ \mu)^* \wedge \phi$.

R^*_1 says that the newly received information is firmly believed. R^*_2 is adapted from R2. We may not have $B \circ \mu = B \cup \{(\mu, 1)\}$ because for formulas in B whose weights are 1, their weights become less than 1 after revision. Other postulates are simply adapted from KM postulates.

We consider the following postulates for iterated revision:

(C1) If $\beta \vdash \mu$, then $((B \circ \mu) \circ \beta)^* \equiv (B \circ \beta)^*$;

(C2) If $\beta \vdash \neg\mu$, then $((B \circ \mu) \circ \beta)^* \equiv (B \circ \beta)^*$;

(C3) If $(B \circ \beta)^* \vdash \mu$ then $(B \circ \mu) \circ \beta \vdash_P \mu$;

(C4) If $(B \circ \beta)^* \wedge \mu$ is consistent, then $((B \circ \mu) \circ \beta)^* \wedge \mu$ is consistent;

(Rec) If $\beta \wedge \mu$ is consistent then $(B \circ \mu) \circ \beta \vdash_P \mu$;

(Ind) If $(B \circ \beta)^* \wedge \mu$ is consistent then $(B \circ \mu) \circ \beta \vdash_P \mu$.

(C1)-(C4) can be found in (Darwiche & Pearl 1997). (Rec) comes from (Nayak, Pagnucco & Peppas 2003) and (Ind) from (Jin & Thielscher 2007).

Proposition 7 Let $\circ$ be a weak-AGM revision operator. When applied to revise possibilistic knowledge base, it satisfies (C3), (C4), (Rec) and (Ind). If $\circ$ satisfies all the KM postulates, then it satisfies (C1). It does not satisfy (C2) in general, even when it satisfies all the KM postulates.

Proposition 7 shows that our operator satisfies most of postulates for iterated revision except (C2), which is controversial according to discussion in (Jin & Thielscher 2007)

Our revision operator also satisfies two interesting postulates for prioritized merging in (Delgrande, Dubois, & Lang 2006).

Proposition 8 Let $\circ$ be a weak-AGM revision operator and μ be a formula. Suppose b_i ($i = 1, ..., l$) are all the distinct weights of formulas appearing in B such that $b_1 > b_2 > ... > b_l > 0$. Let $B_{\geq i} = \{(\phi_j, a_j) \in B : a_j \geq b_i\}$. We then have (1) $B_{\geq i} \circ \mu \vdash_\pi (\phi, a)$ for each $(\phi, a) \in B_{\geq i+1} \circ \mu$; (2) $B_{\geq i} \circ \mu = B_{=i} \circ (B_{\geq i-1} \circ \mu)$, where $B_{=i} = \{(\phi_j, a_j) \in B : a_j = b_i\}$.

Related Work

Iterated revision in possibilistic logic has been discussed in (Benferhat et. al. 2002) where two revision operators are given. However, they have shown that none of their revision operators satisfies DP postulates. Our approach is different from the iterated revision approaches in OCF framework (Darwiche & Pearl 1997; Jin & Thielscher 2007) in that our approach is defined by a one-step revision operator whilst OCF-based revision operator is defined by increasing or decreasing the evidence value of an interpretation to a fixed degree depending on whether it satisfies the new evidence or not.

Our approach extends the revision-based approach in (Qi, Liu & Bell 2005), which generalizes the DMA approach in

(Benferhat et. al. 2004), in two ways. First, the revision-based approach applies a model-based revision operator or a syntax-based revision operator to deal with inconsistency whilst our approach applies a weak-AGM revision operator which is more general than both model-based revision operators and syntax-based revision operators. Second, the revision-based approach returns a set of models when a model-based revision operator is applied whilst our approach semantically returns a possibility distribution. This work is also related to the prioritized merging approach in (Delgrande, Dubois, & Lang 2006) and the model-based merging approach in possibilistic logic in (Qi 2007). Both approaches are based on a one-step (or a flat) merging operator. The prioritized merging approach can be used to deal with iterated revision problem. It satisfies all the KM postulates and some adapted postulates for iterated revision. However, it falsifies most of the original postulates for iterated revision because their approach does not result in a prioritized knowledge base but a flat knowledge base. The model-based merging approach returns a possibility distribution and also gives its syntactical counterpart. However, it cannot be applied to deal with iterated revision. To deal with iterated revision, we apply a weak-AGM revision operator and give a special treatment of the weights of the formulas during revision. It has been shown in (Benferhat et. al. 2004) that DMA approach results in a knowledge base containing more information than those obtained by Adjustment in (Williams 1994) and Maxi-Adjustment in (Williams 1996) when the input for revision is a totally reliable formula. Therefore, our approach is also better than Adjustment and Maxi-Adjustment.

Conclusion and Future Work

In this paper, we have proposed an iterated revision approach in possibilistic logic. We first discussed KM postulates for a one-step revision operator and provided some novel results. We then gave an algorithm for semantic revision in possibilistic logic based on a one-step revision operator. The computation of the semantic approach was given. The algorithm extends the revision-based algorithm in (Qi, Liu & Bell 2005). We showed that our approach can be used to compute some important prioritized base revision operators. Finally, we showed that our revision operator satisfies almost all the DP postulates for iterated revision and other interesting logical properties. We assume that the input formula is fully reliable and has higher priority than the original knowledge. As a future work, we will consider uncertain input formula of the form (ϕ, a), where $a < 1$.

Acknowledgments

This work is partially supported by the EU under the IST project NeOn (IST-2006-027595, `http://www.neon-project.org/`)

References

Benferhat S.; Cayrol, C.; Dubois, D.; Lang, J.; and Prade, H. 1993. Inconsistency Management and Prioritized Syntax-Based Entailment. In *Proc. of IJCAI*, 223-236.

Benferhat S.; Dubois, D.; Prade, H.; and Williams, M.-A. 2002. A practical approach to revising prioritized knowledge bases. *Studia Logica* 70(1):105-130.

Benferhat, S.; Kaci, S.; Berre, D.L.; and Williams, M.-A. 2004. Weakening conflicting information for iterated revision and knowledge integration. *Artif. Intell.* 153(1-2):339-371.

Brewka, G. 1989. Preferred Subtheories: An Extended Logical Framework for Default Reasoning. In *Proc. of IJCAI*, 1043-1048.

Darwiche, A. and Pearl, J. 1997. On the Logic of Iterated Belief Revision. *Artif. Intell.* 89(1-2):1-29.

Delgrande, J.P.; Dubois, D.; and Lang, J. 2006. Iterated revision as prioritized merging. In *Proc. of KR*, 210-220.

Dubois, D.; Lang, J.; and Prade, H. 1994. Possibilistic logic. In Handbook of logic in Artificial Intelligence and Logic Programming, *Volume 3*. Oxford University Press, 439-513.

Fagin, R.; Ullman, J. D.; and Vardi, M. Y. 1983. On the Semantics of Updates in Databases. In *Proc. of PODS*, 352-365.

Gärdenfors, P. 1988. *Knowledge in Flux-Modeling the Dynamic of Epistemic States*. Cambridge: The MIT Press.

Ginsberg, M.L. 1986. Counterfactuals. *Artif. Intell.* 30(1):35-79.

Hansson, S. O. 1999. *A Textbook of Belief Dynamics: Theory Change and Database Updating*. Kluwer Academic Publishers.

Jin, Y. and Thielscher, M. 2007. Iterated belief revision, revised. *Artif. Intell.* 171(1):1-18.

Katsuno, H. and Mendelzon, A.O. 1992. Propositional Knowledge Base Revision and Minimal Change. *Artif. Intell.* 52(3):263-294.

Nayak, A.C.; Pagnucco, M.; and Peppas, P. 2003. Dynamic belief revision operators. *Artif. Intell.* 146(2):193-228.

Nebel, B. 1991. Belief Revision and Default Reasoning: Syntax-Based Approaches. In *Proc. of KR*, 417-428.

Nebel, B. 1998. How Hard is it to Revise a Belief Base? In *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, Kluwer, Dordrecht, The Netherlands, 77-145.

Qi, G.; Liu, W.; and Bell, D. 2005. A Revision-Based Approach to Resolving Conflicting Information. In *Proc. of UAI*, 477-484.

Qi, G. 2007. A Model-based Approach for Merging Prioritized Knowledge Bases in Possibilistic Logic. In *Proc. of AAAI*, 471-476.

Rott, H. 2000. Two dogmas of belief revision. *Journal of Philosophy* 97(9):503-522.

Williams, M.-A. 1994. Transmutations of Knowledge Systems. In *Proc. of KR*, 619-629.

Williams, M.-A. 1996. Towards a Practical Approach to Belief Revision: Reason-Based Change. In *Proc. of KR*, 412-420.