Revising Imprecise Probabilistic Beliefs in the Framework of Probabilistic Logic Programming

Anbu Yue and **Weiru Liu**
School of Electronics, Electrical Engineering and Computer Science,
Queen's University Belfast, Belfast BT7 1NN, UK
{a.yue, w.liu}@qub.ac.uk

Abstract

Probabilistic logic programming is a powerful technique to represent and reason with imprecise probabilistic knowledge. A probabilistic logic program (PLP) is a knowledge base which contains a set of conditional events with probability intervals. In this paper, we investigate the issue of revising such a PLP in light of receiving new information. We propose postulates for revising PLPs when a new piece of evidence is also a probabilistic conditional event. Our postulates lead to Jeffrey's rule and Bayesian conditioning when the original PLP defines a single probability distribution. Furthermore, we prove that our postulates are extensions to Darwiche and Pearl (DP) postulates when new evidence is a propositional formula. We also give the representation theorem for the postulates and provide an instantiation of revision operators satisfying the proposed postulates.

Introduction

Probabilistic logic programming has been used to represent and reason with probabilistic knowledge in many real-world applications, e.g., (Fuhr 2000; De Raedt, Kimmig, and Toivonen 2007; Baral and Hunsaker 2007). Probabilistic knowledge often needs to be revised when new information (evidence) is received. In the literature of probabilistic belief revision and updating, most research so far focuses on revising a single probability distribution (Chan and Darwiche 2005; Grünwald and Halpern 2003; van Fraasen 1976; Dubois and Prade 1997). However, a single probability distribution is not suitable for representing imprecise probabilistic beliefs, as the case for a probabilistic logic program (PLP), where a set of probability distributions is usually associated with a PLP. Research on revising a set of probability distributions is reported in (Skulj 2006; Grove and Halpern 1998), but these methods (as well as methods on revising single probability distributions) can only revise probability distributions by certain kind of evidence, i.e., evidence that is consistent with original distributions. Therefore, any evidence that is not fully consistent with current knowledge cannot be used.

Revision on conditional probabilistic logic is reported in (Kern-Isberner and Rödder 2004). In conditional probabilistic logic, beliefs are represented by a set of probabilistic propositions which can be treated as a set of constraints, and the set of probability distributions that satisfy these constraints define the semantics of the theory. When revising beliefs represented by a set of probabilistic propositions, the revision procedure actually revises a single probability distribution which is obtained by using the maximum entropy principle, therefore some information is lost in the revision procedure (because other distributions are not used). Furthermore, this method also requires that new evidence must be consistent with the original distribution.

On the technical side, most revision methods are based on Bayesian conditioning or its extensions, such as Jeffrey's rule (Skulj 2006) or the minimum cross-entropy principle (van Fraasen 1976; Kern-Isberner and Rödder 2004). A disadvantage of these methods is that a consistent condition is required, so it does not fully satisfy the intension of belief revision in logics. According to the discussion in (Voorbraak 1999), the (extended) conditioning is parallel to belief expansion operators, and should be treated as a type of belief changing operator other than belief revision or updating operators in the logical view. Another disadvantage of conditioning (Bayesian conditioning or Jeffrey's Rule) is that a belief *cannot* be revised by conditional events. Also, applying the minimum cross-entropy principle to revise beliefs by conditional events, some counterintuitive conclusions can be derived (Grove and Halpern 1997).

In this paper, we investigate the issue of revising a PLP in light of receiving new information, which is represented by a probabilistic formula. A revision strategy for revising a PLP shall satisfy the following constraints:

- All knowledge contained by the given PLP should be considered when revised by new evidence. Semantically, the set of probability distributions that satisfy the given PLP and that are closest to new evidence should all be considered in the revision procedure.
- Only elements of knowledge relevant to new evidence should be affected by the revision, that is, *irrelevant* knowledge should be retained. This is also known as the minimal change principle.
- Revision operators should not be too sensitive to the given probability interval of a conditional event (new evidence), that is, by slightly modifying the given probability interval of new evidence, the result of revision should not be

affected significantly.

To achieve these, we first propose postulates to characterize revision operators (strategy) for PLPs, we then present the representation theorem for these postulates and provide an instantiation of revision operators that satisfies our postulates. Our revision strategy has the following properties:

- We do not require that new evidence is consistent with the original PLPs.
- Our revision strategy is an extension of Bayesian conditioning and Jeffrey's rule, in the sense that our postulates lead to Jeffrey's rule and Bayesian conditioning when the original PLP defines a single probability distribution.
- Our revision strategy is also an extension of epistemic revision, in the sense that when new evidence is a sure event with probability 1, our postulates lead to Darwiche and Pearl (DP) postulates for iterated belief revision.

There are other approaches to representing probabilistic knowledge other than PLPs, for example, probabilistic logic programming language P-log. P-log can be used to reason with causal knowledge and new information can be modeled as observations (Baral and Hunsaker 2007). Again, any new information (an observation) must be consistent with the original knowledge.

This paper is organized as follows. In Section 2, we briefly review probabilistic logic programming, DP postulates for iterated belief revision, and introduce the concept of probabilistic epistemic states for PLPs. In Section 3, we propose postulates for revising PLPs with probabilistic conditional events and give the representation theorem to characterize the postulates. Jeffrey's rule and Bayesian conditioning are proved to be subsumed by our revision strategy. In Section 4, we define a specific revision operator that satisfies our postulates. In Section 5, we prove that our postulates extend DP postulates. In Section 6, we discuss related work and in Section 7 we conclude the paper.

Preliminaries

Conditional Probabilistic Logic Programs

We consider conditional probabilistic logic programming in this paper (Lukasiewicz 1998; 2001; 2007).

Let Φ be a finite set of *predicate symbols* and *constant symbols*, and $\mathcal{V}$ be a set of *object variables* and $\mathcal{B}$ be a set of *bound constants* which are in [0,1] describing the bound of probabilities. It is required that Φ contains at least one constant symbol. We use lowercase letters $a, b, \dots$ for constants from Φ, uppercase letters X, Y for object variables, and l, u for bound constants.

An *object term* is a constant from Φ or an object variable from $\mathcal{V}$. An *atom* is of the form $p(t_1, \dots, t_k)$, where p is a predicate symbol and t_i is an object term. An *event* or *formula* is constructed from a set of atoms by logic connectives $\wedge, \vee, \neg$ as usual. We use letters ϕ, ψ, φ for events.

An object term, event, conditional event, probabilistic formula, or PLP is called *ground* iff it does not contain any object variables from $\mathcal{V}$.

Herbrand universe (denoted as HU_Φ) is the set of all constants from Φ, and Herbrand base HB_Φ is the finite nonempty set of all events constructed from the predicate symbols in Φ and constants in HU_Φ. A *possible world* I is a subset of HB_Φ, and $\mathcal{I}_\Phi$ is the set of all possible worlds. An *assignment* σ maps each object variable to an element of HU_Φ. It is extended to object terms by $\sigma(c) = c$ for all constant symbols from Φ. An event ϕ satisfied by I under σ, denoted by $I \models_\sigma \phi$, is defined inductively as:

- $I \models_\sigma p(t_1, \dots, t_n)$ iff $p(\sigma(t_1), \dots, \sigma(t_n)) \in I$;
- $I \models_\sigma \phi_1 \wedge \phi_2$ iff $I \models_\sigma \phi_1$ and $I \models_\sigma \phi_2$;
- $I \models_\sigma \phi_1 \vee \phi_2$ iff $I \models_\sigma \phi_1$ or $I \models_\sigma \phi_2$;
- $I \models_\sigma \neg\phi$ iff $I \not\models_\sigma \phi$

An event φ is satisfied by a possible world I, denoted by $I \models_{cl} \varphi$, iff $I \models_\sigma \varphi$ for all assignments σ. An event φ is a *logical consequence* of event ϕ, denoted as $\phi \models_{cl} \varphi$, iff all possible worlds that satisfy ϕ also satisfy φ.

In this paper, we use $\top$ to represent (ground) tautology, and we have that $I \models_{cl} \top$ for all I and all assignments σ.

A *conditional event* is of the form $\psi|\varphi$ with events ψ and φ. A *probabilistic formula* is of the form $(\psi|\varphi)[l, u]$ which means that the probability bounds for conditional event $\psi|\varphi$ are l and u. We call ψ its *consequent* and φ its *antecedent*. A *probabilistic logic program (PLP)* P is a set of probabilistic formulae.

A *probabilistic interpretation* Pr is a probability distribution on $\mathcal{I}_\Phi$ (i.e., as $\mathcal{I}_\Phi$ is finite, Pr is a mapping from $\mathcal{I}_\Phi$ to the unit interval [0,1] such that $\sum_{I \in \mathcal{I}_\Phi} Pr(I) = 1$). We use $\mathbf{Pr}_\Phi$ to denote the set of all probability distributions on $\mathcal{I}_\Phi$, and use $\mathbf{Pr}$ to denote a subset of $\mathbf{Pr}_\Phi$. The *probability* of an event φ in Pr under an assignment σ, is defined as $Pr_\sigma(\varphi) = \sum_{I \in \mathcal{I}_\Phi, I \models_\sigma \varphi} Pr(I)$. If φ is ground, we simply write as $Pr(\varphi)$. If $Pr_\sigma(\varphi) > 0$ then we define $Pr_\sigma(\psi|\varphi) = Pr_\sigma(\varphi \wedge \psi)/Pr_\sigma(\varphi)$, and $Pr_{\sigma,\varphi}(I) = Pr(I)/Pr_\sigma(\varphi)$ if $I \models_\sigma \varphi$ or $Pr_{\sigma,\varphi}(I) = 0$ otherwise. When φ is ground, we simply write $Pr_{\sigma,\varphi}(I)$ as $Pr_\varphi(I)$.

A probabilistic formula $(\psi|\varphi)[l, u]$ is satisfied by a probabilistic interpretation Pr under an assignment σ, denoted by: $Pr \models_\sigma (\psi|\varphi)[l, u]$ iff $Pr_\sigma(\varphi) = 0$ or $Pr_\sigma(\psi|\varphi) \in [l, u]$. A probabilistic formula $(\psi|\varphi)[l, u]$ is *satisfied* by a probabilistic interpretation Pr, or Pr is a *probabilistic model* of $(\psi|\varphi)[l, u]$, denoted by $Pr \models (\psi|\varphi)[l, u]$, iff $Pr \models_\sigma (\psi|\varphi)[l, u]$ for all assignments σ. A probabilistic interpretation is a *probabilistic model* of a PLP P, denoted by $Pr \models P$, iff Pr is a probabilistic model of all $(\psi|\varphi)[l, u] \in P$. A PLP P is *satisfiable* or *consistent* iff a model of P exists. A probabilistic formula $(\psi|\varphi)[l, u]$ is a *consequence* of the PLP P, denoted by $P \models (\psi|\varphi)[l, u]$, iff all probabilistic models of P are also probabilistic models of $(\psi|\varphi)[l, u]$. A probabilistic formula $(\psi|\varphi)[l, u]$ is a *tight consequence* of P, denoted by $P \models_{tight} (\psi|\varphi)[l, u]$, iff $P \models (\psi|\varphi)[l, u]$, $P \not\models (\psi|\varphi)[l, u']$, $P \not\models (\psi|\varphi)[l', u]$ for all $l' > l$ and $u' < u$ ($l', u' \in [0, 1]$). Note that, if $P \models (\phi|\top)[0, 0]$, then it is canonically defined as $P \models_{tight} (\psi|\phi)[1, 0]$, where $[1, 0]$ stands for an empty set.

Iterated Belief Revision on Epistemic States

Belief revision is a process of changing a belief set to accommodate new evidence that is possibly inconsistent with existing beliefs, and the process is regulated by the

AGM (Alchourrón, Gärdenfors, and Markinson) postulates (Alchourrón, Gärdenfors, and Makinson 1985; Gärdenfors 1998), and rephrased by Katsuno and Mendelzon for propositional logic setting (Katsuno and Mendelzon 1992). Darwiche and Pearl (Darwiche and Pearl 1997) modified these postulates based on epistemic states and proposed additional postulates for iterated belief revision. Each epistemic state, Ψ_{pl}, has a unique belief set, denoted as $Bel_{pl}(\Psi_{pl})$, which can be taken as a propositional sentence. It is possible to have two different epistemic states with the equivalent belief set. When an epistemic state is embedded in a logical formula, it is used to represent its belief set. For example, $\Psi_{pl} \models \alpha$ means $Bel_{pl}(\Psi_{pl}) \models \alpha$, $\Psi_{pl} \wedge \phi$ stands for $Bel_{pl}(\Psi_{pl}) \wedge \phi$, and $\Psi_{pl} \equiv \Psi'_{pl}$ stands for $Bel_{pl}(\Psi_{pl}) \equiv Bel_{pl}(\Psi'_{pl})$.

Let $\circ$ be a revision operator, the modified AGM postulates for epistemic revision are (note: $\Psi_{pl} \circ \beta$ is an epistemic state) (Darwiche and Pearl 1997):

G1 $\Psi_{pl} \circ \beta$ implies β.
G2 If $\Psi_{pl} \wedge \beta$ is satisfiable, then $\Psi_{pl} \circ \beta \equiv \Psi_{pl} \wedge \beta$.
G3 If β is satisfiable, then $\Psi_{pl} \circ \beta$ is also satisfiable.
G4 If $\Psi_{pl} = \Phi_{pl}$ and $\beta_1 \equiv \beta_2$, then $\Psi_{pl} \circ \beta_1 \equiv \Phi_{pl} \circ \beta_2$.
G5 $(\Psi_{pl} \circ \beta) \wedge \phi$ implies $\Psi_{pl} \circ (\beta \wedge \phi)$.
G6 If $(\Psi_{pl} \circ \beta) \wedge \phi$ is satisfiable, then $\Psi \circ (\beta \wedge \phi)$ implies $(\Psi_{pl} \circ \beta) \wedge \phi$.

and the postulates for iterated belief revision are (note: $\Psi_{pl} \models (\beta|\alpha)$ if $\Psi_{pl} \circ \alpha \models \beta$)

C1 If $\alpha \models \phi$, then $\Psi_{pl} \models (\beta|\alpha)$ iff $\Psi_{pl} \circ \phi \models (\beta|\alpha)$.
C2 If $\alpha \models \neg\phi$, then $\Psi_{pl} \models (\beta|\alpha)$ iff $\Psi_{pl} \circ \phi \models (\beta|\alpha)$.
C3 If $\Psi_{pl} \models (\beta|\alpha)$ then $\Psi_{pl} \circ \beta \models (\beta|\alpha)$.
C4 If $\Psi_{pl} \not\models (\neg\beta|\alpha)$ then $\Psi_{pl} \circ \beta \not\models (\neg\beta|\alpha)$.

In this paper, we refer the above G1 to G6 and C1 to C4 postulates as DP postulates.

Probabilistic Epistemic States and Belief Set

Given a PLP P, we define set $Bel^0(P)$ as $Bel^0(P) = \{(\psi|\phi)[l,u] \mid P \models (\psi|\phi)[l,u], P \not\models (\phi|\top)[0,0]\}$ and call it the belief set of P. Condition $P \not\models (\phi|\top)[0,0]$ is required because when $P \models (\phi|\top)[0,0]$, $P \models (\psi|\phi)[l,u]$ for all ψ and all $[l,u] \subseteq [0,1]$. Without this condition, some counterintuitive conclusions can be inferred, for instance, $(\psi|\phi)[0,0.3]$ and $(\psi|\phi)[0.9,1]$ can simultaneously be the beliefs of an agent if $P \models (\phi|\top)[0,0]$.

Each *probabilistic epistemic state*, Ψ, has a unique belief set, denoted as $Bel^0(\Psi)$, which is a set of probabilistic formulae. $Bel^0(\Psi)$ is *closed*, i.e. $Bel^0(Bel^0(\Psi)) = Bel^0(\Psi)$. We call Ψ a probabilistic epistemic state of a PLP P, iff $Bel^0(\Psi) = Bel^0(P)$. It is possible to have more than one probabilistic epistemic state for the same PLP. On the other hand, each *probabilistic epistemic state* Ψ also has a unique set of probability distributions, denoted $Mods(\Psi)$. It is required that $Mods(\Psi)$ is a convex set. $Mods(\Psi)$ defines the semantics of $Bel^0(\Psi)$, that is, $Bel^0(\Psi) = \{(\psi|\phi)[l,u] \mid \exists Pr \in Mods(\Psi), Pr(\phi) > 0$, and $\forall Pr \in Mods(\Psi)$, if $Pr(\phi) > 0$, then $Pr(\psi|\phi) \in [l,u]\}$. We call the set $Bel_{cl}(\Psi) = \{(\phi|\top)[1,1] | (\phi|\top)[1,1] \in Bel^0(\Psi)\}$ as the *classical belief set of* Ψ.

Furthermore, we have the following inference relations:
$\Psi \models (\psi|\phi)[l,u]$ iff $(\psi|\phi)[l,u] \in Bel^0(\Psi)$, and
$\Psi \models_{tight} (\psi|\phi)[l,u]$ iff $\Psi \models (\psi|\phi)[l,u]$ and for all $[l',u'] \subset [l,u]$, $\Psi \not\models (\psi|\phi)[l,u]$. It is worth noting that, when $\Psi \models (\phi|\top)[0,0]$, $\Psi \not\models_{tight} (\psi|\phi)[l,u]$ for all $l,u \in [0,1]$, however $P \models_{tight} (\psi|\phi)[1,0]$.

In this paper, we also write $\Psi \wedge (\psi|\phi)[l,u]$ to represent $Bel^0(\Psi) \cup \{(\psi|\phi)[l,u]\}$.

Imprecise Probabilistic Belief Revision

Postulates

Definition 1 (perpendicular) *A conditional event $(\psi|\phi)$ is perpendicular with another conditional event $(\psi'|\phi')$, denoted as $(\psi|\phi) \bowtie (\psi'|\phi')$ iff at least one of the following conditions holds*

- $\phi \models_{cl} \phi' \wedge \psi'$,
- $\phi \models_{cl} \phi' \wedge \neg\psi'$,
- $\models_{cl} \neg(\phi' \wedge \phi)$,
- $\phi' \models_{cl} \phi \wedge \psi$,
- $\phi' \models_{cl} \phi \wedge \neg\psi$,

The perpendicularity relation formalizes a kind of *irrelevance* between two conditional events. The above definition is an extension of the definition of perpendicular in (Kern-Isberner 1999), in which the first three conditions are required. Conditional event $(\psi'|\phi')$ only affects the relationship (probability distributions) between $\phi' \wedge \psi'$ and $\phi' \wedge \neg\psi'$, and has no effect on sub-events of $\phi' \wedge \psi'$ and of $\phi' \wedge \neg\psi'$. When the first condition holds, $(\psi'|\phi')$ shall have no effect on $(\psi|\phi)$ since the latter is a sub-event of ϕ and hence a sub-event of $\phi' \wedge \psi'$. Therefore, events $(\psi|\phi)$ and $(\psi'|\phi')$ are irrelevant. The 2nd, 4th and the 5th conditions can be explained similarly. When the third condition is true, events $(\psi|\phi)$ and $(\psi'|\phi')$ affect different domains of events, so they are irrelevant.

Definition 2 *Let P be a PLP with probabilistic epistemic state Ψ and $\mu = (\psi|\phi)[l,u]$ be a probabilistic formula. The result of revising Ψ by μ is another probabilistic epistemic state, denoted as $\Psi \star \mu$ where $\star$ is a revision operator. Operator $\star$ is required to satisfy the following postulates:*

R1 $\Psi \star \mu \models \mu$.
R2 $\Psi \wedge \mu \models \Psi \star \mu$.
R3 *if $\Psi \wedge \mu$ is satisfiable, then $\Psi \star \mu \models \Psi \wedge \mu$.*
R4 *$\Psi \star \mu$ is unsatisfiable only if μ is unsatisfiable.*
R5 *$\Psi \star \mu \equiv \Psi \star \mu'$ if $\mu \equiv \mu'$.*
R6 *Let $\mu = (\psi|\phi)[l,u]$ and $\Psi \star \mu \models_{tight} (\psi|\phi)[l',u']$. Let $\mu' = (\psi|\phi)[l_1,u_1]$ and $\Psi \star \mu' \models_{tight} (\psi|\phi)[l'_1,u'_1]$. For any $\epsilon > 0$, if $|u_1 - u| + |l_1 - l| < \epsilon$, and both of $(\psi|\phi)[l,u]$ and $(\psi|\phi)[l',u']$ are satisfiable, then $|u'_1 - u'| + |l'_1 - l'| < \epsilon$.*
R7 *Let $\mu = (\psi|\phi)[l,u]$, if $\Psi \models (\phi|\top)[l',u']$, then $(\Psi \star \mu) \models (\phi|\top)[l',u']$.*
R8 *Let $\mu = (\psi|\phi)[l,u]$ and $\mu' = (\psi'|\phi')[l',u']$. Suppose that $(\psi|\phi) \bowtie (\psi'|\phi')$. If $(\Psi \star \mu) \wedge \mu'$ is satisfiable then $\Psi \wedge \mu'$ is satisfiable, and $(\Psi \star \mu) \wedge \mu' = (\Psi \wedge \mu') \star \mu$.*

R1 - R5 is an analog to postulates G1 - G4. We do not have corresponding postulates for G5 and G6 since revision with the conjunction of probabilistic formulae are more complicated and is beyond the scope of this paper. R6 is a *sensitivity* requirement, which says that a slightly modification on the bounds of $\mu = (\psi|\phi)[l, u]$ (i.e., $\mu' = (\psi|\phi)[l_1, u_1]$) shall not affect the result of revision significantly. R7 says that revising Ψ by $\mu = (\psi|\phi)[l, u]$ should not affect the statement about ϕ (, but the impreciseness of ϕ may be decreased). Recall that perpendicular condition characterizes a kind of irrelevance, R8 says that any irrelevance knowledge with new evidence should not be lost by the revision with this evidence. It is worth noting that R7 is a special case of R8, since $(\phi|\top) \bowtie (\psi|\phi)$ for all ψ.

From R7, it is possible to have $\Psi \models_{tight} (\phi|\top)[l, u]$, $\Psi \star \mu \models_{tight} (\phi|\top)[l', u']$ and $[l', u'] \subset [l, u]$, which means that the ignorance about ϕ is reduced after revision. To reflect this, a stronger version of R7 to totally retain the ignorance about ϕ can be defined as

R7' Let $\mu = (\psi|\phi)[l, u]$, $\Psi \models (\phi|\top)[l', u']$ iff $(\Psi \star \mu) \models (\phi|\top)[l', u']$.

However, R7' is too strong and it is inconsistent with R3. Let us see the following example. Suppose that

$$P = \{(fly(t)|bird(t))[0.98, 1], (bird(t)|penguin(t))[1, 1]\}.$$

Let Ψ be the probabilistic epistemic state of P, then $\Psi \models_{tight} (penguin(t)|\top)[0, 1]$. If a revision operator $\star$ satisfies R7', then $\Psi \star (fly(t)|penguin(t))[0, 0] \models_{tight} (penguin(t)|\top)[0, 1]$. However, $\Psi \wedge (fly(t)|penguin(t))[0, 0] \models (penguin(t)|\top)[0, 0.02]$, so $\star$ does not satisfy R3.

From R1, it is clear that $\Psi \star \mu$ may infer a tighter bound than $[l, u]$ for conditional event $(\psi|\phi)$. This leads to the following proposition.

Proposition 1 *Let Ψ be a probabilistic epistemic state, $\star$ be a revision operator, and $\mu = (\psi|\phi)[l, u]$ be a probabilistic formula. Suppose that $\Psi \models_{tight} (\psi|\phi)[l', u']$. If $\star$ satisfies R1, R2, R3, and R6 then*

- $\Psi \star \mu \models (\psi|\phi)[l, l]$ *if* $l \geq u'$,
- $\Psi \star \mu \models (\psi|\phi)[u, u]$ *if* $u \leq l'$,
- $\Psi \star \mu \models (\psi|\phi)[l_1, u_1]$, *where* $l_1 = \max\{l, l'\}$ *and* $u_1 = \min\{u, u'\}$, *otherwise.*

That is, after the revision, we can infer a tighter bound for the conditional event $(\psi|\phi)$ and this tighter bound is either $[l, l]$, or $[u, u]$, or $[l', u']$. This bound, among all the sub-bounds of $[l, u]$ that is given in the new evidence, is also the *closest* to the original beliefs.

Definition 3 *Let ϕ be an event, and Pr_1, Pr_2 be two probability distributions. We can define an event ψ s.t. $I \models_{cl} \psi$ iff $Pr_1(I) \geq Pr_2(I)$. We define a function d_ϕ that maps a pair of probability distributions to a non-negative real number, s.t. $d_\phi(Pr_1, Pr_2) = \max\{d_p * Pr_2(\psi|\phi), d_n * Pr_2(\neg\psi|\phi)\}$, where*

$$d_p = d_n = \infty + 1, \text{ if } \exists I \models_{cl} \neg\phi, Pr_1(I) \neq Pr_2(I)$$

otherwise

$$\begin{cases} d_p = \max_{I \models_{cl} \phi} \frac{Pr_1(I) - Pr_2(I)}{Pr_2(I)}, I \models_{cl} \psi \\ d_n = \max_{I \models_{cl} \phi} \frac{Pr_2(I) - Pr_1(I)}{Pr_2(I)}, I \models_{cl} \neg\psi \end{cases}$$

Here, we define $\frac{0}{0} = 0$ and $\frac{a}{0} = \infty, 0 * \infty = 0, a > 0$.

From d_ϕ, we can define function $dis_\phi : \mathbf{Pr}_\Phi \times 2^{\mathbf{Pr}_\Phi} \mapsto R^+ \cup \{0\}$ by $dis_\phi(Pr, \mathbf{Pr}) = \min_{Pr' \in \mathbf{Pr}} d_\phi(Pr, Pr')$

In this paper, we define $\infty + 1 > \infty$.

Obviously, $d_\phi(Pr_1, Pr_2) \geq 0$, and d_ϕ is asymmetric. d_ϕ gives a quasi-distance from one distribution to another, and dis_ϕ characterizes how close a probability distribution to a set of probability distributions is.

Definition 4 *Let Ψ be a probabilistic epistemic state and ϕ be an event. Suppose that $\Psi \models_{tight} (\phi|\top)[l, u]$. A function that maps each epistemic state Ψ to a total pre-order $\leq_{\phi,\Psi}$ on the set of all possible probabilistic distributions $\mathbf{Pr}_\Phi$ is said to be a* **faithful assignment** *if and only if*

1. *if both Pr_1 and Pr_2 are in $Mods(\Psi)$ then $Pr_1 =_{\phi,\Psi} Pr_2$*
2. *if $Pr_1(\phi) \in [l, u]$, $Pr_2(\phi) \notin [l, u]$, then $Pr_1 <_{\phi,\Psi} Pr_2$*
3. *if $Pr_1(\phi) \in [l, u]$ and $Pr_2(\phi) \in [l, u]$, then $Pr_1 <_{\phi,\Psi} Pr_2$ whenever $dis_\phi(Pr_1, Mods(\Psi)) < dis_\phi(Pr_2, Mods(\Psi))$*

For any probabilistic epistemic state Ψ, and any event ϕ, if $\leq_{\phi,\Psi}$ is mapped with Ψ by a faithful assignment, then $Mods(\Psi) = \min(\mathbf{Pr}_\Phi, \leq_{\phi,\Psi})$, where $\min(\mathbf{Pr}_\Phi, \leq_{\phi,\Psi})$ contains all the probability distributions that are minimal in $\mathbf{Pr}_\Phi$ according to the total pre-order $\leq_{\phi,\Psi}$.

Theorem 1 (Representation Theorem) *A revision operator $\star$ satisfies postulates R1-R8 precisely when there exists a faithful assignment that maps each epistemic state Ψ and event ϕ to a total pre-order $\leq_{\phi,\Psi}$ such that*

$$Mods(\Psi \star (\psi|\phi)[l, u]) = \min(Mods((\psi|\phi)[l, u]), \leq_{\phi,\Psi})$$

Here $Mods((\psi|\phi)[l, u])$ is defined as $Mods((\psi|\phi)[l, u]) = \{Pr | Pr \models (\psi|\phi)[l, u], Pr \in \mathbf{Pr}_\Phi\}$.

Ramsey Test

Traditionally, a conditional $(B|A)$ (meaning *if A then B*) is interpreted by the Ramsey test as follows:

Ramsey test $\Sigma \models (B|A)$ iff $\Sigma \star A \models B$ where Σ is a belief set and $\star$ is a revision operator.

The Ramsey test says that a conditional is accepted iff when revising the beliefs with the antecedent (e.g., A, so that the antecedent is true), the revised beliefs must also make the consequent (e.g. B) true. We adapt the Ramsey test to the probabilistic logic setting.

Probabilistic Ramsey test: if $\Psi \models (\psi|\phi)[l, u]$ then $\Psi \star (\phi|\top)[1, 1] \models (\psi|\top)[l, u]$, and if $\Psi \star (\phi|\top)[1, 1] \models (\psi|\top)[l, u]$, $\Psi \not\models (\phi|\top)[0, 0]$, then $\Psi \models (\psi|\phi)[l', u']$ with $l' \leq l, u' \geq u$.

Proposition 2 *Let Ψ be a probabilistic epistemic state and $\star$ be a revision operator. If $\star$ satisfies R8, then $\star$ satisfies the probabilistic Ramsey test.*

A Specific Revision Operator

In this section, we define a specific probabilistic epistemic state Ψ_P of a PLP P as a convex set of probability distributions, and we define $Mods(\Psi_P) = \Psi_P$. Therefore, $Bel^0(\Psi) = \{(\psi|\phi)[l,u] \mid \exists Pr \in \Psi_P, Pr(\phi) > 0$, and $\forall Pr \in \Psi_P$, if $Pr(\phi) > 0$, then $Pr(\psi|\phi) \in [l,u]\}$. Obviously, if Ψ_P is a probabilistic epistemic state of PLP P and $\mathbf{Pr} = \{Pr \mid Pr \models P\}$, then $\Psi_P \subseteq \mathbf{Pr}$.

Example 1 *Let Ψ be a set of probability distributions and $\Psi = \{Pr \mid Pr(p_1(t)) + Pr(p_2(t)) = 0.8\}$, where p_1 and p_2 are two predicates and t is a constant symbol. Let P be a PLP such that*

$$P = \left\{ \begin{array}{ll} (p_1(t)|\top)[0, 0.8], & (p_1(t) \vee p_2(t)|\top)[0.4, 0.8], \\ (p_2(t)|\top)[0, 0.8], & (p_1(t) \wedge p_2(t)|\top)[0, 0.4] \end{array} \right\}$$

It is easy to check that Ψ is a convex set and $Bel^0(\Psi) = Bel^0(P)$, therefore, Ψ is a probabilistic epistemic state of P. Suppose that $\mathbf{Pr} = \{Pr \mid Pr \models P\}$. Obviously, $\Psi \subset \mathbf{Pr}$, which means that Ψ provides more information than $\mathbf{Pr}$. Trivially, $\mathbf{Pr}$ is also a probabilistic epistemic state of P.

Definition 5 *Let Pr be a probability distribution and $\mu = (\psi|\phi)[l,u]$ be a probabilistic formula. We define function dis as $dis(Pr,\mu) = \min_{x\in[l,u]} |a - x|$ that maps a Pr and μ to a non negative real number where $Pr(\psi|\phi) = a$. If $Pr(\phi) = 0$ then we define $dis(Pr,\mu) = 0$.*

Let Ψ_P be the probabilistic epistemic state of PLP P defined above and $\mu = (\psi|\phi)[l,u]$ be a probabilistic formula. Suppose that $Pr \in \Psi_P$, and $dis(Pr,\mu) = \min_{Pr'\in\Psi_P} dis(Pr',\mu)$. If $dis(Pr,\mu) = 0$, then $Pr \models \mu$. If $dis(Pr,\mu) > 0$, then $Pr(\phi) > 0$ and $Pr(\psi|\phi)$ is the closest probability value to the bound $[l,u]$ given in μ. Therefore, Pr is one of the probability distributions that is the closest to probabilistic formula μ. The following specific revision operator is defined based on this principle, that is, only those probability distributions in Ψ_P that are closest to new evidence (μ) should be selected and revised.

Definition 6 *Let Ψ_P be the probabilistic epistemic state of a PLP P, and $\mu = (\psi|\phi)[l,u]$ be a probabilistic formula. We define an operator $\bullet$ such that $\Psi'_P = \Psi_P \bullet \mu$ is another set of probability distributions such that $Pr' \in \Psi'_P$ iff there exists $Pr \in \Psi_P$ s.t. $dis(Pr,\mu) = \min_{Pr_0\in\Psi_P} dis(Pr_0,\mu)$, and*

- *$Pr' = Pr$, if $dis(Pr,\mu) = 0$, or*
- *$dis(Pr,\mu) > 0$, $Pr(\psi|\phi) = b$, and*
  - *$Pr'(I) = Pr(I)$, for all $I \models_{cl} \neg\phi$*
  - *if $b > 0$, then $Pr'(I) = \frac{a}{b}Pr(I)$ for all $I \models_{cl} \phi \wedge \psi$, if $b = 0$, then $\sum_{I\models_{cl}\phi\wedge\psi} Pr'(I) = a * Pr(\phi)$,*
  - *if $b < 1$, then $Pr'(I) = \frac{1-a}{1-b}Pr(I)$ for all $I \models_{cl} \phi \wedge \neg\psi$, if $b = 1$, then $\sum_{I\models_{cl}\phi\wedge\neg\psi} Pr'(I) = (1-a) * Pr(\phi)$,*

  where $a = \arg\min_{x\in[l,u]} |x - b|$

Assume that $Pr \in \Psi_P$, $dis(Pr,\mu) = \min_{Pr_0\in\Psi_P} dis(Pr_0,\mu)$, and Pr' is obtained from Pr based on the above definition. If $dis(Pr,\mu) > 0$, then the values of b and a can be calculated, and $Pr'(\psi|\phi) = a$. We can also prove that, for all $Pr'' \in \mathbf{Pr}_\Phi$, $d_\phi(Pr', Pr) \leq d_\phi(Pr'', Pr)$ if $Pr''(\psi|\phi) = a$.

Proposition 3 *The set $\Psi'_P = \Psi_P \bullet \mu$ is a convex set.*

Proposition 4 *The revision operator $\bullet$ satisfies postulates R1-R8.*

When the agent's belief is precise, i.e., $Bel^0(\Psi) \models_{tight} (\psi|\phi)[a,a]$ for all ψ and ϕ such that $Bel^0(\Psi) \not\models (\phi|\top)[0,0]$, there is one and only one probability distribution that satisfying $Bel^0(\Psi)$, in another word, $Mods(\Psi)$ is a singleton set. On the other hand, if $Mods(\Psi)$ is a singleton set, then $Bel^0(\Psi)$ is precise. In this case, revising Ψ collapses to Jeffrey's rule and Bayesian conditioning.

Theorem 2 *Let Ψ_P be a probabilistic epistemic state such that $Mods(\Psi_P)$ is a singleton set and ϕ be an event. Suppose that $\Psi'_P = \Psi_P \bullet (\phi|\top)[l,l]$, then Ψ'_P is a singleton set. Assume that $\Psi_P = \{Pr\}$ and $\Psi'_P = \{Pr'\}$,*

- *If $Pr(\phi) > 0$ and $l = 1$, then $Pr'(I) = Pr_\phi(I)$, that is, Pr' is obtained from Pr by Bayesian conditioning on ϕ.*
- *If $0 < Pr(\phi) < 1$, then $Pr'(I) = l * Pr_\phi(I) + (1 - l) * Pr_{\neg\phi}(I)$, that is, Pr' is obtained from Pr by Jeffrey's rule.*

Example 2 *Let PLP P be defined as follows:*

$$P = \left\{ \begin{array}{l} (fly(X)|bird(X))[0.98, 1] \\ (bird(X)|penguin(X))[1, 1] \\ (penguin(X)|bird(X))[0.1, 1] \end{array} \right\}$$

Table 1: Probability distributions in Ψ_P and $\Psi_P \bullet \mu$

$Index$	$Possible\ world$	$Pr \in \Psi_P$	$Pr' \in \Psi_P \bullet \mu$
I_1	$\emptyset$	x_1	x'_1
I_2	$\{penguin(t)\}$	ϵ_1	ϵ'_1
I_3	$\{bird(t)\}$	x_2	x'_2
I_4	$\{bird(t), penguin(t)\}$	x_3	x'_3
I_5	$\{fly(t)\}$	x_4	x'_4
I_6	$\{fly(t), penguin(t)\}$	ϵ_2	ϵ'_2
I_7	$\{fly(t), bird(t)\}$	x_5	x'_5
I_8	$\{fly(t), bird(t), penguin(t)\}$	x_6	x'_6

where

$x_1 + x_2 + x_3 + x_4 + x_5 + x_6 + \epsilon_1 + \epsilon_2 = 1$,

$\frac{x_5+x_6}{x_2+x_3+x_5+x_6} \in [0.98, 1]$ or $x_2 + x_3 + x_5 + x_6 = 0$,

$\frac{x_3+x_6}{x_3+x_6+\epsilon_1+\epsilon_2} = 1$ or $x_3 + x_6 + \epsilon_1 + \epsilon_2 = 0$,

$\frac{x_3+x_6}{x_2+x_3+x_5+x_6} \in [0.1, 1]$ or $x_2 + x_3 + x_5 + x_6 = 0$

thus $\epsilon_1 = \epsilon_2 = 0$

Then we have $P \models_{tight} (fly(t)|penguin(t))[0.8, 1]$. When new evidence suggests that penguins cannot fly, we revise our knowledge with evidence $\mu = (fly(X)|penguin(X))[0, 0]$.

Suppose that $\Psi_P = \{Pr \mid Pr \models P\}$. The probability distributions in Ψ_P that are closest to μ are those with $Pr(fly(t)|penguin(t)) = 0.8$, $Pr(penguin(t)|\top) = 0.1$, and $Pr(fly(t)|bird(t)) = 0.98$. Therefore, we revise the subset of Ψ_P such that each probability distribution in the subset having $x_3 + x_6 = 0.1 * (x_2 + x_3 + x_5 + x_6)$, $x_6 = 0.8*(x_3+x_6)$, and $x_5+x_6 = 0.98*(x_2+x_3+x_5+x_6)$. By simplification, the above constraints can be re-expressed as $x_1+x_4 = 1-a, x_2 = 0, x_3 = 0.02*a, x_5 = 0.9*a, x_6 = 0.08*a, \epsilon_1 = \epsilon_2 = 0$, where $a \in [0, 1]$.

So, $x_3 = 0.02*(x_2+x_3+x_5+x_6), x_6 = 0.08*(x_2+x_3+x_5+x_6)$. According to μ, we do not need to revise the models that do not satisfy $penguin(t)$, i.e. $x'_1 = x_1, x'_2 = x_2, x'_4 = x_4, x'_5 = x_5$. From the definition of $\bullet$, we also have that $\epsilon'_1 = \epsilon_1, \epsilon'_2 = \epsilon_2 = 0$, and $x'_6 = \frac{0}{0.8} * x_6 = 0$, $x'_3 = \frac{1-0}{1-0.8} * x_3 = 0.1*(x_2+x_3+x_5+x_6)$. Therefore, a probability distribution in $\Psi_P \bullet \mu$ must satisfy that $x'_1+x'_4 = 1-a, x'_2 = 0, x'_3 = 0.1*a, x'_5 = 0.9*a, x'_6 = 0, \epsilon'_1 = \epsilon'_2 = 0$.

Let $\Psi'_P = \Psi_P \bullet \mu$ be the revised probabilistic epistemic state, then

$$\Psi'_P \models (fly(t)|bird(t))[0.9, 0.9]$$
$$\Psi'_P \models (bird(t)|penguin(t))[1, 1]$$
$$\Psi'_P \models (penguin(t)|bird(t))[0.1, 0.1]$$

We can also infer that $(penguin(X)|bird(X))[0.1, 0.1]$, which means that the revised knowledge about the proportions of penguins in birds is getting more precise. Because only when the proportion of penguins (in birds) is small, we can have **the probability of a bird can fly** *as close to* 1 *as possible. We also have $\Psi'_P \models (fly(t)|bird(t))[0.9, 0.9]$ which means that definitely not all birds can fly given the new evidence.*

Relationship with Darwiche-Pearl Revision

In this section, we prove that when a probabilistic epistemic state (from a PLP) is revised by a sure formula (rather than a probabilistic formula), we can induce C1-C4 in (Darwiche and Pearl 1997) on iterated revision, and therefore our probabilistic revision strategy is an extension of their work.

Proposition 5 *Let P be a PLP, Ψ be its epistemic state, and $\star$ be a revision operator that satisfies postulates R1 - R8. Suppose that ϕ, α, β are events, then $\star$ also satisfies the following postulates:*

PC1 *If $\alpha \models_{cl} \phi$, then $\Psi \models (\beta|\alpha)[l, u]$ iff $\Psi \star (\phi|\top)[l_1, u_1] \models (\beta|\alpha)[l, u]$.*
Explanation: *Accommodating evidence about ϕ should not perturb any conditional beliefs that are conditioned on a premise more specific than ϕ.*

PC2 *If $\Psi \models_{tight} (\beta|\alpha)[l, u]$, then $\Psi \star (\beta|\top)[l_1, u_1] \models_{tight} (\beta|\alpha)[l', u']$, where $l' \geq l$ and $u' \geq u$.*
Explanation: *The lower bound and upper bound of $(\beta|\alpha)$ should not be decreased after accommodating evidence that support β.*

From postulate PC1, we can get the following instantiations:

- If $\alpha \models_{cl} \phi$, then $\Psi \models (\beta|\alpha)[l, u]$ iff $\Psi \star (\phi|\top)[1, 1] \models (\beta|\alpha)[l, u]$
- If $\alpha \models_{cl} \neg\phi$, then $\Psi \models (\beta|\alpha)[l, u]$ iff $\Psi \star (\neg\phi|\top)[1, 1] \models (\beta|\alpha)[l, u]$

It is easy to see that the above two instantiations are extensions of the DP postulates C1 and C2 respectively.

Intuitively, $\Psi_{pl} \not\models (\neg\beta|\alpha)$ means that it is impossible to infer $(\neg\beta|\alpha)$. In probabilistic logic programs, this is modeled by $\Psi \not\models (\neg\beta|\alpha)[l, u]$ for all $l > 0$ and it can be equivalently rewritten as $\Psi \models_{tight} (\beta|\alpha)[l', 1]$, for some $l' \in [0, 1]$.

Similarly, we can get instantiations of PC2 as:

- If $\Psi \models_{tight} (\beta|\alpha)[1, 1]$, then $\Psi \star (\beta|\top)[1, 1] \models_{tight} (\beta|\alpha)[1, 1]$.
- If $\Psi \models_{tight} (\beta|\alpha)[l, 1]$, then $\Psi \star (\beta|\top)[1, 1] \models_{tight} (\beta|\alpha)[l', 1]$, where $l' \geq l$.

It is easy to see that the above two instantiations are extensions of the DP postulates C3 and C4 respectively.

Related Work

In (Grove and Halpern 1998), the authors axiomatized the process of updating a set of probability measures, and stated that *it is unwise to simply assume that a result in the standard model (single-measure model) can be trivially* **lifted** *to apply in general model (set-measure model).* Some postulates were provided to justify the process of updating a set of probability distributions, and some updating operators were examined according to these postulates. By analyzing these postulates in our revision framework, we find that their postulate P3 and P5 are too strong.

P3 $(\mathbf{Pr} \star_\Phi \phi) \star_\Phi \phi' = \mathbf{Pr} \star_\Phi (\phi \wedge \phi')$

P5 $\mathbf{Pr} \star_\Phi \phi = \bigcup_{Pr \in \mathbf{Pr}} \{Pr\} \star_\Phi \phi$

If an operator satisfies P3 and R4 defined in our paper, then the consistency of new evidence with the original beliefs is required. In fact, P3 is meaningful only when $\phi \wedge \phi'$ is consistent.

P5 requires that all probability distributions equally contribute to the revision process. This requirement is too strong. Let $\mathbf{Pr} = \{Pr \mid Pr \models P\}$, and $\mathbf{Pr}' = \mathbf{Pr} \star_\Phi \phi$. If P is consistent with $(\phi|\top)[1, 1]$, then $\mathbf{Pr}' \subseteq \mathbf{Pr}$ by postulate **R3**. Suppose that $Pr \in \mathbf{Pr}$ and $Pr(\phi) \neq 1$. Then $Pr \notin \mathbf{Pr}'$ but $\{Pr\} \star_\Phi \phi \in \mathbf{Pr}'$. Since P is arbitrarily selected, it can be inferred that $\{Pr\} \star_\Phi \phi = \emptyset$ or $\mathbf{Pr}' = \{Pr \mid Pr(\phi) = 1\}$. If $\{Pr\} \star_\Phi \phi = \emptyset$, then $\star_\Phi$ does not satisfy postulates **R2** and **R4**. On the other hand, $\mathbf{Pr}' = \{Pr \mid Pr(\phi) = 1\}$ means that the revision process is irrelevant to the original $\mathbf{Pr}$, this is counter-intuitive. Therefore, a revision operator should not satisfy **R2**, **R3**, **R4** and **P5** simultaneously.

Since R2, R3, and R4 are widely accepted postulates (from modified AGM postulates) for revision, P3 and P5 (Grove and Halpern 1998) are not appropriate postulates for revising imprecise probabilistic beliefs in the logical view. Our revision strategy satisfies the rest of the postulates, e.g, P1, P2, P4, P6*, and P7.

In (Kern-Isberner 1999), postulates (CR0-CR7) were provided for revising epistemic states by conditional beliefs.

In contract, we proposed postulates for revising imprecise probabilistic beliefs. In (Kern-Isberner 2002), it was proved that using an operator (defined from the minimal cross entropy principle) to revise a single probability distribution by a single conditional event satisfies postulates CR5 - CR7 in (Kern-Isberner 1999). However, revising by applying minimal cross-entropy principle does not satisfy our R7 or R8.

In (Voorbraak 1999), the authors studied conditioning and constraining in the logical view. Constraining is the process of selecting the probability distributions that give probability 1 to the new evidence. The author argued that constraining was comparable to expansion but conditioning was different from revision in logical view. To some extend, the specific operator defined in this paper is a kind of combination of conditioning and constraining, however, we proved that our operator satisfies DP postulates.

Revision on propositional probabilistic logic was reported in (Kern-Isberner and Rödder 2004). One disadvantage is that new evidence is required to be Pr-consistent with respect to the original probability distribution Pr. The revision operator in (Kern-Isberner and Rödder 2004) considers only one distribution obtained by maximum entropy, so the knowledge contained by other distributions are ignored in the revision process. In contract, our (specific) revision operator works on a set of probabilistic distributions. Also, their operator does not satisfy postulates our R7 or R8.

In (Grove and Halpern 1997), the authors stated that any revision or updating should work on a *second-order* probability distribution, which is a distribution in the space of combination of possible distributions and events of knowing the conditional probability of a conditional event. However, our method does not require this second-order distribution and therefore is arguably simpler.

Conclusion

In this paper, we proposed postulates for imprecise probabilistic belief revision, which extend modified AGM postulates and DP postulates. Therefore, logically PLPs can be taken as a formal language to represent belief sets of probabilistic epistemic states. In traditional (iterated) belief revision, new evidence is generally a propositional formula, however, in our revision framework new evidence can be a probabilistic formula. Furthermore, our postulates are proved to be extensions of Jeffrey's rule and Bayesian conditioning, when a PLP defines a single probability distribution.

In traditional probabilistic logic, belief revising and updating are restricted to requiring new evidence to be consistent with the original beliefs. In practice, such requirement is too strong to be satisfied in general. On the contrary, our revision framework does not require this condition on new evidence.

References

Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *J. Symb. Log.* 50(2):510–530.

Baral, C., and Hunsaker, M. 2007. Using the probabilistic logic programming language p-log for causal and counterfactual reasoning and non-naive conditioning. In *Proc. of IJCAI'07*, 243–249.

Chan, H., and Darwiche, A. 2005. On the revision of probabilistic beliefs using uncertain evidence. *Artif. Intell.* 163(1):67–90.

Darwiche, A., and Pearl, J. 1997. On the logic of iterated belief revision. *Artif. Intell.* 89(1-2):1–29.

De Raedt, L.; Kimmig, A.; and Toivonen, H. 2007. Problog: A probabilistic prolog and its application in link discovery. In *Proc. IJCAI'07*, 2462–2467.

Dubois, D., and Prade, H. 1997. Focusing vs. belief revision: A fundamental distinction when dealing with generic knowledge. In *Proc. of ECSQARU-FAPR'97*, 96–107.

Fuhr, N. 2000. Probabilistic datalog: Implementing logical information retrieval for advanced applications. *JASIS* 51(2):95–110.

Gärdenfors, P. 1998. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge, Mass.

Grove, A. J., and Halpern, J. Y. 1997. Probability update: Conditioning vs. cross-entropy. In *Proc. of UAI'97*, 208–214.

Grove, A. J., and Halpern, J. Y. 1998. Updating sets of probabilities. In *Proc. of UAI'98*, 173–182.

Grünwald, P., and Halpern, J. Y. 2003. Updating probabilities. *J. Artif. Intell. Res. (JAIR)* 19:243–278.

Katsuno, H., and Mendelzon, A. O. 1992. Propositional knowledge base revision and minimal change. *Artif. Intell.* 52(3):263–294.

Kern-Isberner, G., and Rödder, W. 2004. Belief revision and information fusion on optimum entropy. *Int. J. Intell. Syst.* 19(9):837–857.

Kern-Isberner, G. 1999. Postulates for conditional belief revision. In *Proc. of IJCAI'99*, 186–191.

Kern-Isberner, G. 2002. The principle of conditional preservation in belief revision. In *Proc. of FoIKS'02*, 105–129.

Lukasiewicz, T. 1998. Probabilistic logic programming. In *Proc. of ECAI'98*, 388–392.

Lukasiewicz, T. 2001. Probabilistic logic programming with conditional constraints. *ACM Trans. Comput. Log.* 2(3):289–339.

Lukasiewicz, T. 2007. Nonmonotonic probabilistic logics under variable-strength inheritance with overriding: Complexity, algorithms, and implementation. *Int. J. Approx. Reasoning* 44(3):301–321.

Skulj, D. 2006. Jeffrey's conditioning rule in neighbourhood models. *Int. J. Approx. Reasoning* 42(3):192–211.

van Fraasen, B. 1976. Probabilities of conditionals. In *Proc. of Foundations of Probability Theory, Statistical Inference, and Statistical Theories of Science*, 261–300.

Voorbraak, F. 1999. Probabilistic belief change: Expansion, conditioning and constraining. In *Proc. of UAI'99*, 655–662.