

Weak Nonmonotonic Probabilistic Logics

Thomas Lukasiewicz*

Dipartimento di Informatica e Sistemistica,
Università degli Studi di Roma “La Sapienza”
Via Salaria 113, I-00198 Rome, Italy
lukasiewicz@dis.uniroma1.it

Abstract

Towards probabilistic formalisms for resolving local inconsistencies under model-theoretic probabilistic entailment, we present probabilistic generalizations of Pearl’s entailment in System Z and Lehmann’s lexicographic entailment. We then analyze the nonmonotonic and semantic properties of the new notions of entailment. In particular, we show that they satisfy the rationality postulates of System P and the property of Rational Monotonicity. Moreover, we show that model-theoretic probabilistic entailment is stronger than the new notion of lexicographic entailment, which in turn is stronger than the new notion of entailment in System Z . As an important feature of the new notions of entailment in System Z and lexicographic entailment, we show that they coincide with model-theoretic probabilistic entailment whenever there are no local inconsistencies. We also show that the new notions of entailment in System Z and lexicographic entailment are proper generalizations of their classical counterparts. Finally, we present algorithms for reasoning under the new formalisms, and we give a precise picture of its computational complexity.

Introduction

During the recent decades, reasoning about probabilities has started to play an important role in AI. In particular, reasoning about interval restrictions for conditional probabilities, also called conditional constraints (which are of the form $(\psi|\phi)[l, u]$ with a conditional event $\psi|\phi$ and reals $l, u \in [0, 1]$) has been a subject of extensive research efforts.

One important approach for handling conditional constraints is model-theoretic probabilistic logic, which has its origin in philosophy and logic, and whose roots can be traced back to Boole (1854). There is a wide spectrum of formal languages that have been explored in model-theoretic probabilistic logic, ranging from constraints for unconditional and conditional events to rich languages that specify linear inequalities over events (see especially the works by Nilsson (1986), Fagin *et al.* (1990), Dubois and Prade *et al.* (1988; 1991), Frisch and Haddawy (1994), and Lukasiewicz (1999a; 1999b; 2001b); see also the survey on sentential

probability logic by Hailperin (1996)). The main decision and optimization problems in model-theoretic probabilistic logic are deciding satisfiability, deciding logical consequence, and computing tight logically entailed intervals.

The notion of model-theoretic probabilistic entailment is widely accepted in AI. However, it fails to produce satisfactory conclusions about conditional events $\psi|\phi$ in the case where our knowledge about the premise ϕ is locally inconsistent. The following example illustrates this drawback.

Example 1 Suppose that we have the knowledge “all penguins are birds”, “birds have legs with probability 1”, “birds fly with probability 1”, and “penguins fly with a probability of at most 0.05”. Under model-theoretic probabilistic logic, this knowledge is *locally inconsistent relative to penguins* (since it says “penguins fly with probability 1”, but in the same time also “penguins fly with a probability of at most 0.05”), and thus we cannot conclude anything about the properties of penguins. In fact, we even conclude “there are no penguins”, which reports this local inconsistency related to penguins. But it seems reasonable to interpret “birds have legs with probability 1” and “birds fly with probability 1” as defaults “generally, birds have legs” and “generally, birds fly”, respectively, and to conclude “generally, penguins have legs” (that is, “penguins have legs with probability 1”) and “penguins fly with a probability of at most 0.05”. □

The above example suggests that such local inconsistencies under model-theoretic probabilistic entailment are essentially due to the fact that 0/1-probabilistic knowledge is interpreted as strict logical knowledge, and not as default knowledge. The main idea behind this paper is to combine the notion of model-theoretic probabilistic entailment with mechanisms from default reasoning from conditional knowledge bases in order to obtain a notion of entailment in probabilistic logic that can handle such local inconsistencies.

The literature contains several different proposals for default reasoning from conditional knowledge bases and extensive work on its desired properties. The core of these properties are the rationality postulates of System P by Kraus, Lehmann, and Magidor (1990), which constitute a sound and complete axiom system for several classical model-theoretic entailment relations under uncertainty measures on worlds. They characterize classical model-theoretic entailment under preferential structures (Shoham 1987;

*Alternate address: Institut für Informationssysteme, Technische Universität Wien, Favoritenstraße 9-11, A-1040 Vienna, Austria; e-mail: lukasiewicz@kr.tuwien.ac.at.
Copyright © 2004, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

Kraus, Lehmann, & Magidor 1990), infinitesimal probabilities (Adams 1975; Pearl 1989), possibility measures (Dubois & Prade 1991), and world rankings (Spohn 1988; Goldszmidt & Pearl 1992). They also characterize an entailment relation based on conditional objects (Dubois & Prade 1994). A survey of all these relationships is given in (Benferhat, Dubois, & Prade 1997; Gabbay & Smets 1998). Mainly to solve problems with irrelevant information, the notion of rational closure as a more adventurous notion of entailment was introduced by Lehmann (1989; 1992). It is equivalent to entailment in System Z by Pearl (1990), to the least specific possibility entailment by Benferhat *et al.* (1992), and to a conditional (modal) logic-based entailment by Lamarre (1992). Finally, mainly to solve problems with property inheritance from classes to exceptional subclasses, lexicographic entailment was introduced by Lehmann (1995) and Benferhat *et al.* (1993), and other more sophisticated notions of entailment for default reasoning were proposed.

In this paper, we present probabilistic generalizations of Pearl's entailment in System Z (1990) and of Lehmann's lexicographic entailment (1995), which are weaker than entailment in model-theoretic probabilistic logic. Roughly, model-theoretic probabilistic entailment realizes an inheritance of 0/1-probabilistic knowledge along subclass relationships, but does not have any mechanism for resolving local inconsistencies due to this inheritance. The new probabilistic formalisms now add such a mechanism to model-theoretic probabilistic entailment, and this is why they are weaker than model-theoretic probabilistic entailment. The main contributions of this paper are as follows:

- We present probabilistic generalizations of Pearl's entailment in System Z and Lehmann's lexicographic entailment, which are weaker than model-theoretic probabilistic entailment. We explore and compare the nonmonotonic properties of the new notions of entailment and of model-theoretic probabilistic entailment. In particular, the new formalisms satisfy the rationality postulates of System P , the property of Rational Monotonicity, and some Irrelevance and Direct Inference properties.
- We analyze the relationship between the new notion of entailment in System Z , the new notion of lexicographic entailment, and model-theoretic probabilistic entailment. It turns out that model-theoretic probabilistic entailment is stronger than the new lexicographic entailment, which is in turn stronger than the new entailment in System Z .
- As an important feature of the new entailment in System Z and the new lexicographic entailment, we show that *they coincide with model-theoretic probabilistic entailment whenever there are no local inconsistencies* as illustrated in Example 1. Furthermore, we show that they are proper generalizations of their classical counterparts.
- We present algorithms for computing tight intervals under the new notions of entailment in System Z and of lexicographic entailment, which are based on reductions to the standard tasks of deciding model existence and computing tight intervals under model-theoretic probabilistic entailment. Furthermore, we give a precise picture of the

complexity of deciding logical consequence and of computing tight intervals under the new notions of entailment in System Z and of lexicographic entailment in general as well as restricted cases.

Note that detailed proofs of all results in this paper are given in the extended paper (Lukasiewicz 2003).

Model-Theoretic Probabilistic Logic

In this section, we recall the main concepts from model-theoretic probabilistic logic (see especially the works by Nilsson (1986), Fagin *et al.* (1990), Dubois and Prade *et al.* (1988; 1991), Frisch and Haddawy (1994), and Lukasiewicz (1999a; 1999b; 2001b).

We assume a set of *basic events* $\Phi = \{p_1, \dots, p_n\}$, $n \geq 1$, and use \perp and \top to denote *false* and *true*, respectively. The set of *events* is inductively defined as follows. Every element of $\Phi \cup \{\perp, \top\}$ is an event. If ϕ and ψ are events, then also $\neg\phi$ and $(\phi \wedge \psi)$. A *conditional event* has the form $\psi|\phi$, where ψ and ϕ are events. A *conditional constraint* is of the form $(\psi|\phi)[l, u]$ with a conditional event $\psi|\phi$ and reals $l, u \in [0, 1]$. We define *probabilistic formulas* inductively as follows. Every conditional constraint is a probabilistic formula. If F and G are probabilistic formulas, then also $\neg F$ and $(F \wedge G)$. We use $(F \vee G)$ (resp., $(F \Leftarrow G)$) to abbreviate $\neg(\neg F \wedge \neg G)$ (resp., $\neg(\neg F \wedge G)$), where F and G are either two events or two probabilistic formulas, and adopt the usual conventions to eliminate parentheses. A *logical constraint* is an event of the form $\psi \Leftarrow \phi$. A *probabilistic knowledge base* $KB = (L, P)$ consists of a finite set of logical constraints L and a finite set of conditional constraints P .

A *world* I is a truth assignment to the basic events in Φ , which is inductively extended to all events as usual (that is, by $I(\perp) = \text{false}$, $I(\top) = \text{true}$, $I(\neg\phi) = \text{true}$ iff $I(\phi) = \text{false}$, and $I(\phi \wedge \psi) = \text{true}$ iff $I(\phi) = I(\psi) = \text{true}$). We use \mathcal{I}_Φ to denote the set of all worlds for Φ . A world I is a *model* of an event ϕ , denoted $I \models \phi$, iff $I(\phi) = \text{true}$. We say I is a *model* of a set of events L , denoted $I \models L$, iff I is a model of all $\phi \in L$. An event ϕ (resp., a set of events L) is *satisfiable* iff a model of ϕ (resp., L) exists. An event ψ is a *logical consequence* of ϕ (resp., L), denoted $\phi \models \psi$ (resp., $L \models \psi$), iff each model of ϕ (resp., L) is also a model of ψ . We use $\phi \not\models \psi$ (resp., $L \not\models \psi$) to denote that $\phi \models \psi$ (resp., $L \models \psi$) does not hold.

A *probabilistic interpretation* Pr is a probability function on \mathcal{I}_Φ (that is, a mapping Pr from \mathcal{I}_Φ to $[0, 1]$ such that $\sum_{I \in \mathcal{I}_\Phi} Pr(I) = 1$). The *probability* of an event ϕ in Pr is defined as $Pr(\phi) = \sum_{I \in \mathcal{I}_\Phi, I \models \phi} Pr(I)$. For events ϕ and ψ with $Pr(\phi) > 0$, let $Pr(\psi|\phi) = Pr(\psi \wedge \phi) / Pr(\phi)$, and let the *conditioning* of Pr on ϕ be defined by $Pr_\phi(I) = Pr(I) / Pr(\phi)$ for all $I \in \mathcal{I}_\Phi$ with $I \models \phi$, and by $Pr_\phi(I) = 0$ for all other $I \in \mathcal{I}_\Phi$. The *truth* of logical constraints and probabilistic formulas F in Pr , denoted $Pr \models F$, is defined by induction as follows:

- $Pr \models \psi \Leftarrow \phi$ iff $Pr(\psi \wedge \phi) = Pr(\phi)$;
- $Pr \models (\psi|\phi)[l, u]$ iff $Pr(\phi) = 0$ or $Pr(\psi|\phi) \in [l, u]$;
- $Pr \models \neg F$ iff not $Pr \models F$; and
- $Pr \models (F \wedge G)$ iff $Pr \models F$ and $Pr \models G$.

We say Pr satisfies F , or Pr is a *model* of F , iff $Pr \models F$. It satisfies a set of logical constraints and probabilistic formulas \mathcal{F} , or Pr is a *model* of \mathcal{F} , denoted $Pr \models \mathcal{F}$, iff Pr is a model of all $F \in \mathcal{F}$. We say \mathcal{F} is *satisfiable* iff a model of \mathcal{F} exists. A logical constraint or probabilistic formula F is a *logical consequence* of \mathcal{F} , denoted $\mathcal{F} \models F$, iff every model of \mathcal{F} is also a model of F .

A probabilistic knowledge base $KB = (L, P)$ is *satisfiable* iff $L \cup P$ is satisfiable. We next define the notion of *logical entailment* for conditional constraints from KB . Note that each entailment relation for conditional constraints consists of a consequence relation and a tight consequence relation. A conditional constraint $(\psi|\phi)[l, u]$ is a *logical consequence* of KB , denoted $KB \models (\psi|\phi)[l, u]$, iff $L \cup P \models (\psi|\phi)[l, u]$. We say $(\psi|\phi)[l, u]$ is a *tight logical consequence* of KB , denoted $KB \models_{tight} (\psi|\phi)[l, u]$, iff l (resp., u) is the infimum (resp., supremum) of $Pr(\psi|\phi)$ subject to all models Pr of $L \cup P$ with $Pr(\phi) > 0$. Note that we define $[l, u]$ as the empty interval $[1, 0]$, when $L \cup P \models \perp \Leftarrow \phi$.

Example 2 The knowledge that “all eagles are birds”, “birds have legs with the probability 1”, and “birds fly with a probability of at least 0.95” can be expressed by the probabilistic knowledge base KB_1 shown in Table 1.

In model-theoretic probabilistic logic, KB_1 encodes the *strict logical knowledge* “all eagles are birds” and “all birds have legs” (that is, in model-theoretic probabilistic logic, a logical constraint $\psi \Leftarrow \phi \in L$ has the same meaning as a conditional constraint $(\psi|\phi)[1, 1] \in P$), and the *probabilistic knowledge* “birds fly with a probability of at least 0.95”.

It is not difficult to see that KB_1 is satisfiable, and that some tight logical consequences of KB_1 are given as shown in Table 2. Notice that the 0/1-probabilistic property of having legs is inherited from birds to the subclass eagles, while the probabilistic property of being able to fly with a probability of at least 0.95 is *not* inherited from birds to eagles. \square

Example 3 The knowledge “all penguins are birds”, “birds have legs with the probability 1”, “birds fly with the probability 1”, and “penguins fly with a probability of at most 0.05” can be expressed by the probabilistic knowledge base $KB_2 = (L_2, P_2)$ shown in Table 1. It is not difficult to see that KB_2 is satisfiable, and that some tight logical consequences of KB_2 are as shown in Table 2.

Here, the empty interval “[1, 0]” for the last two conditional events is due to the fact that the 0/1-probabilistic property of being able to fly is inherited from birds to penguins and is incompatible there with penguins being able to fly with a probability of at most 0.05. That is, our knowledge about penguins is inconsistent. That is, there does not exist any model Pr of $L_2 \cup P_2$ such that $Pr(penguin) > 0$, and thus we are having a local inconsistency relative to *penguin*. Hence, logical entailment is too strong here, since the desired tight conclusions from KB_2 are $(fly|penguin)[0, 0.05]$ and $(legs|penguin)[1, 1]$ instead of $(fly|penguin)[1, 0]$ and $(legs|penguin)[1, 0]$, respectively. \square

Weak Nonmonotonic Probabilistic Logics

In this section, we present novel probabilistic generalizations of Pearl’s entailment in System Z and of Lehmann’s

Table 1: Probabilistic Knowledge Bases

$KB_1 = (\{bird \Leftarrow eagle\}, \{(legs bird)[1, 1], (fly bird)[0.95, 1]\})$
$KB_2 = (\{bird \Leftarrow penguin\}, \{(legs bird)[1, 1], (fly bird)[1, 1], (fly penguin)[0, 0.05]\})$

Table 2: Tight Conclusions

KB	$(\psi \phi)$	\models_{tight}	\sim_{tight}^{lex}	\sim_{tight}^z	\sim_{tight}^p
KB_1	$(legs bird)$	$[1, 1]$	$[1, 1]$	$[1, 1]$	$[1, 1]$
KB_1	$(fly bird)$	$[0.95, 1]$	$[0.95, 1]$	$[0.95, 1]$	$[0.95, 1]$
KB_1	$(legs eagle)$	$[1, 1]$	$[1, 1]$	$[1, 1]$	$[0, 1]$
KB_1	$(fly eagle)$	$[0, 1]$	$[0, 1]$	$[0, 1]$	$[0, 1]$
KB_2	$(legs bird)$	$[1, 1]$	$[1, 1]$	$[1, 1]$	$[1, 1]$
KB_2	$(fly bird)$	$[1, 1]$	$[1, 1]$	$[1, 1]$	$[1, 1]$
KB_2	$(legs penguin)$	$[1, 0]$	$[1, 1]$	$[0, 1]$	$[0, 1]$
KB_2	$(fly penguin)$	$[1, 0]$	$[0, 0.05]$	$[0, 0.05]$	$[0, 0.05]$

lexicographic entailment. We first define probability rankings, and a notion of entailment that is based on sets of probability rankings, which generalizes entailment in System P , and which coincides with probabilistic entailment under g-coherence (see below). We then define the novel formalisms, which are based on unique single probability rankings.

Example 4 Under weak nonmonotonic probabilistic logics, KB_1 in Table 1 represents the *strict logical knowledge* “all eagles are birds”, the *default knowledge* “generally, birds have legs” (that is, a logical constraint $\psi \Leftarrow \phi \in L$ now *does not have* anymore the same meaning as a conditional constraint $(\psi|\phi)[1, 1] \in P$; note that only $(\psi|\phi)[0, 0]$ and $(\psi|\phi)[1, 1]$ in P express defaults), and the *probabilistic knowledge* “birds fly with a probability of at least 0.95”. \square

Preliminaries

A probabilistic interpretation Pr verifies a conditional constraint $(\psi|\phi)[l, u]$ iff $Pr(\phi) > 0$ and $Pr \models (\psi|\phi)[l, u]$. We say Pr falsifies $(\psi|\phi)[l, u]$ iff $Pr(\phi) > 0$ and $Pr \not\models (\psi|\phi)[l, u]$. A set of conditional constraints P tolerates a conditional constraint C under a set of logical constraints L iff $L \cup P$ has a model that verifies C . We say P is *under L in conflict* with C iff no model of $L \cup P$ verifies C .

A *conditional constraint ranking* σ on a probabilistic knowledge base $KB = (L, P)$ maps each $C \in P$ to a non-negative integer. It is *admissible* with KB iff every $P' \subseteq P$ that is under L in conflict with some $C \in P$ contains some C' such that $\sigma(C') < \sigma(C)$.

In the sequel, we use $\alpha > 0$ to abbreviate the probabilistic formula $\neg(\alpha|\top)[0, 0]$. A *probability ranking* κ maps each probabilistic interpretation on \mathcal{I}_Φ to a member of $\{0, 1, \dots\} \cup \{\infty\}$ such that $\kappa(Pr) = 0$ for at least one interpretation Pr . It is extended to all logical constraints and probabilistic formulas F as follows. If F is satisfiable,

then $\kappa(F) = \min \{ \kappa(Pr) \mid Pr \models F \}$; otherwise, $\kappa(F) = \infty$. A probability ranking κ is *admissible* with a probabilistic knowledge base $KB = (L, P)$ iff $\kappa(\neg F) = \infty$ for all $F \in L$ and $\kappa(\phi > 0) < \infty$ and $\kappa(\phi > 0 \wedge (\psi|\phi)[l, u]) < \kappa(\phi > 0 \wedge \neg(\psi|\phi)[l, u])$ for all $(\psi|\phi)[l, u] \in P$.

Consistency and Entailment in System P

We now generalize the notions of consistency and entailment in System P to probabilistic knowledge bases.

A probabilistic knowledge base $KB = (L, P)$ is *p-consistent* iff there exists a probability ranking κ that is admissible with KB . We then define the notion of *p-entailment* in terms of admissible probability rankings as follows. A conditional constraint $(\psi|\phi)[l, u]$ is a *p-consequence* of a *p-consistent* $KB = (L, P)$, denoted $KB \models^p (\psi|\phi)[l, u]$, iff $\kappa(\phi > 0) = \infty$ or $\kappa(\phi > 0 \wedge (\psi|\phi)[l, u]) < \kappa(\phi > 0 \wedge \neg(\psi|\phi)[l, u])$ for every probability ranking κ admissible with KB . We say that $(\psi|\phi)[l, u]$ is a *tight p-consequence* of KB , denoted $KB \models_{tight}^p (\psi|\phi)[l, u]$, iff $l = \sup l'$ (resp., $u = \inf u'$) subject to $KB \models^p (\psi|\phi)[l', u']$.

In ordinary default reasoning, the notion of *p-consistency* is equivalent to the existence of admissible default rankings (Geffner 1992). The following theorem shows that similarly probabilistic *p-consistency* can be expressed in terms of admissible conditional constraint rankings.

Theorem 5 *A probabilistic knowledge base $KB = (L, P)$ is p-consistent iff there exists a conditional constraint ranking on KB that is admissible with KB .*

The next theorem shows that also a characterization of ordinary *p-consistency* due to Goldszmidt and Pearl (1991) carries over to probabilistic *p-consistency*.

Theorem 6 *A probabilistic knowledge base $KB = (L, P)$ is p-consistent iff there is an ordered partition (P_0, \dots, P_k) of P such that either (a) every P_i , $0 \leq i \leq k$, is the set of all $C \in \bigcup_{j=i}^k P_j$ tolerated under L by $\bigcup_{j=i}^k P_j$, or (b) for every i , $0 \leq i \leq k$, each $C \in P_i$ is tolerated under L by $\bigcup_{j=i}^k P_j$.*

The following result shows that also a characterization of ordinary *p-entailment*, which is essentially due to Adams (1975), carries over to the probabilistic case.

Theorem 7 *Let $KB = (L, P)$ be a p-consistent probabilistic knowledge base and $(\beta|\alpha)[l, u]$ be a conditional constraint. Then, $KB \models^p (\beta|\alpha)[l, u]$ iff $(L, P \cup \{(\beta|\alpha)[p, p]\})$ is not p-consistent for all $p \in [0, l) \cup (u, 1]$.*

The next result completes the picture.

Theorem 8 *Let $KB = (L, P)$ be a p-consistent probabilistic knowledge base, and let $(\beta|\alpha)[l, u]$ be a conditional constraint. Then, $KB \models_{tight}^p (\beta|\alpha)[l, u]$ iff (i) $(L, P \cup \{(\beta|\alpha)[p, p]\})$ is not p-consistent for all $p \in [0, l) \cup (u, 1]$, and (ii) $(L, P \cup \{(\beta|\alpha)[p, p]\})$ is p-consistent for all $p \in [l, u]$.*

It is easy to verify that the probabilistic knowledge bases KB_1 and KB_2 in Table 1 are both *p-consistent*. Some tight conclusions under *p-entailment* are shown in Table 2. Observe that neither the default property of having legs, nor the probabilistic property of being able to fly with a probability of at least 0.95, is inherited from birds down to eagles.

Entailment in System Z

We now extend entailment in System Z (Pearl 1990; Goldszmidt & Pearl 1996) to *p-consistent* probabilistic knowledge bases $KB = (L, P)$. The new notion of entailment in System Z is associated with an ordered partition of P , a conditional constraint ranking z on KB , and a probability ranking κ^z . The *z-partition* of KB is the unique ordered partition (P_0, \dots, P_k) of P such that each P_i is the set of all $C \in \bigcup_{j=i}^k P_j$ that are tolerated under L by $\bigcup_{j=i}^k P_j$.

Example 9 The *z-partition* of KB_1 in Table 1 is given by

$$(P_0) = (\{(legs \mid bird)[1, 1], (fly \mid bird)[0.95, 1]\}),$$

while the *z-partition* of KB_2 in Table 1 is given by

$$(P_0, P_1) = (\{(legs \mid bird)[1, 1], (fly \mid bird)[1, 1]\}, \{(fly \mid penguin)[0, 0.05]\}). \quad \square$$

We next define z and κ^z . For every $j \in \{0, \dots, k\}$, each $C \in P_j$ is assigned the value j under the conditional constraint ranking z . The probability ranking κ^z on all probabilistic interpretations Pr is then defined by:

$$\kappa^z(Pr) = \begin{cases} \infty & \text{if } Pr \not\models L \\ 0 & \text{if } Pr \models L \cup P \\ 1 + \max_{C \in P: Pr \not\models C} z(C) & \text{otherwise.} \end{cases}$$

The following lemma shows that the rankings z and κ^z are both admissible with KB .

Lemma 10 *Let $KB = (L, P)$ be p-consistent. Then, z and κ^z are both admissible with KB .*

We next define a preference relation on probabilistic interpretations as follows. For probabilistic interpretations Pr and Pr' , we say Pr is *z-preferable* to Pr' iff $\kappa^z(Pr) < \kappa^z(Pr')$. A model Pr of a set of logical constraints and probabilistic formulas \mathcal{F} is a *z-minimal model* of \mathcal{F} iff no model of \mathcal{F} is *z-preferable* to Pr .

We finally define the notion of *z-entailment* as follows. A conditional constraint $(\psi|\phi)[l, u]$ is a *z-consequence* of KB , denoted $KB \models^z (\psi|\phi)[l, u]$, iff every *z-minimal model* of $L \cup \{\phi > 0\}$ satisfies $(\psi|\phi)[l, u]$. We say $(\psi|\phi)[l, u]$ is a *tight z-consequence* of KB , denoted $KB \models_{tight}^z (\psi|\phi)[l, u]$, iff l (resp., u) is the infimum (resp., supremum) of $Pr(\psi|\phi)$ subject to all *z-minimal models* Pr of $L \cup \{\phi > 0\}$.

Example 11 Table 2 gives the tight conclusions under *z-entailment* from the probabilistic knowledge bases in Table 1. They show that *z-entailment* realizes an inheritance of 0/1-probabilistic properties from classes to non-exceptional subclasses. But it does not inherit 0/1-probabilistic properties from classes to subclasses that are exceptional relative to some other property (and thus, like its classical counterpart, has the problem of inheritance blocking). \square

The following theorem characterizes the notion of *z-consequence* in terms of the probability ranking κ^z .

Theorem 12 *Let $KB = (L, P)$ be a p-consistent probabilistic knowledge base, and let $(\psi|\phi)[l, u]$ be a conditional constraint. Then, $KB \models^z (\psi|\phi)[l, u]$ iff $\kappa^z(\phi > 0) = \infty$ or $\kappa^z(\phi > 0 \wedge (\psi|\phi)[l, u]) < \kappa^z(\phi > 0 \wedge \neg(\psi|\phi)[l, u])$.*

Lexicographic Entailment

We next extend Lehmann's lexicographic entailment (1995) to p -consistent probabilistic knowledge bases $KB = (L, P)$. Note that, even though we do not use probability rankings here, the new notion of lexicographic entailment can be easily expressed through a unique single probability ranking.

We use the z -partition (P_0, \dots, P_k) of KB to define a lexicographic preference relation on probabilistic interpretations as follows. For probabilistic interpretations Pr and Pr' , we say that Pr is *lexicographically preferable* (or *lex-preferable*) to Pr' iff some $i \in \{0, \dots, k\}$ exists such that $|\{C \in P_i \mid Pr \models C\}| > |\{C \in P_i \mid Pr' \models C\}|$ and $|\{C \in P_j \mid Pr \models C\}| = |\{C \in P_j \mid Pr' \models C\}|$ for all $i < j \leq k$. A model Pr of a set of logical constraints and probabilistic formulas \mathcal{F} is a *lexicographically minimal* (or *lex-minimal*) model of \mathcal{F} iff no model of \mathcal{F} is *lex-preferable* to Pr .

We are now ready to define the notion of *lexicographic entailment* (or *lex-entailment*) as follows. A conditional constraint $(\psi|\phi)[l, u]$ is a *lex-consequence* of KB , denoted $KB \models^{lex} (\psi|\phi)[l, u]$, iff each *lex-minimal* model of $L \cup \{\phi > 0\}$ satisfies $(\psi|\phi)[l, u]$. We say $(\psi|\phi)[l, u]$ is a *tight lex-consequence* of KB , denoted $KB \models^{tight} (\psi|\phi)[l, u]$, iff l (resp., u) is the infimum (resp., supremum) of $Pr(\psi|\phi)$ subject to all *lex-minimal* models Pr of $L \cup \{\phi > 0\}$.

Example 13 Table 2 gives the tight conclusions under *lex-entailment* from KB_1 and KB_2 in Table 1. They show that *lex-entailment* realizes a correct inheritance of logical properties, without the problem of inheritance blocking. \square

Semantic Properties

In this section, we explore the semantic properties of the probabilistic notions of p -, z -, and *lex-entailment*, and we give a comparison to logical entailment. We first describe their nonmonotonicity and nonmonotonic properties. We then explore the relationships between the probabilistic formalisms and to their classical counterparts.

Nonmonotonicity

Logical entailment has the following property of *inheritance of logical knowledge (L-INH)* along subclass relationships:

L-INH. If $KB \models (\psi|\phi)[c, c]$ and $\phi \Leftarrow \phi^*$ is valid, then $KB \models (\psi|\phi^*)[c, c]$,

for all events ψ , ϕ , and ϕ^* , all probabilistic knowledge bases KB , and all $c \in \{0, 1\}$. The notions of p -, z -, and *lex-entailment* are nonmonotonic in the sense that they all do not satisfy *L-INH*. Here, p -entailment completely fails *L-INH*, while z - and *lex-entailment* realize some weaker form of *L-INH*, as they are both obtained from logical entailment by adding some strategy for resolving local inconsistencies.

Note that logical, p -, z -, and *lex-entailment* all do not have the following property of *inheritance of purely probabilistic knowledge (P-INH)* along subclass relationships:

P-INH. If $KB \models (\psi|\phi)[l, u]$ and $\phi \Leftarrow \phi^*$ is valid, then $KB \models (\psi|\phi^*)[l, u]$,

for all events ψ , ϕ , and ϕ^* , all probabilistic knowledge bases KB , and all $[l, u] \subseteq [0, 1]$ different from $[0, 0]$, $[1, 1]$, and

$[1, 0]$. See (Lukasiewicz 2002) for entailment semantics that satisfy *P-INH* and restricted forms of *P-INH*. For example, under such entailment semantics, we can conclude $(fly | eagle)[0.95, 1]$ from KB_1 in Table 1.

Nonmonotonic Properties

We now explore the nonmonotonic behavior (especially related to *L-INH*) of the probabilistic formalisms of this paper.

We first consider the postulates *Right Weakening (RW)*, *Reflexivity (Ref)*, *Left Logical Equivalence (LLE)*, *Cut*, *Cautious Monotonicity (CM)*, and *Or* proposed by Kraus, Lehmann, and Magidor (1990), which are commonly regarded as being particularly desirable for any reasonable notion of nonmonotonic entailment. The following result shows that the notions of logical, p -, z -, and *lex-entailment* all satisfy (probabilistic versions of) these postulates.

Theorem 14 \models , \models^p , \models^z , and \models^{lex} satisfy the following properties for all probabilistic knowledge bases $KB = (L, P)$, all events ε , ε' , ϕ , and ψ , and all $l, l', u, u' \in [0, 1]$:

RW. If $(\phi \top)[l, u] \Rightarrow (\psi \top)[l', u']$ is logically valid and $KB \models (\phi|\varepsilon)[l, u]$, then $KB \models (\psi|\varepsilon)[l', u']$.

Ref. $KB \models (\varepsilon|\varepsilon)[1, 1]$.

LLE. If $\varepsilon \Leftrightarrow \varepsilon'$ is logically valid, then $KB \models (\phi|\varepsilon)[l, u]$ iff $KB \models (\phi|\varepsilon')[l, u]$.

Cut. If $KB \models (\varepsilon|\varepsilon')[1, 1]$ and $KB \models (\phi|\varepsilon \wedge \varepsilon')[l, u]$, then $KB \models (\phi|\varepsilon')[l, u]$.

CM. If $KB \models (\varepsilon|\varepsilon')[1, 1]$ and $KB \models (\phi|\varepsilon')[l, u]$, then $KB \models (\phi|\varepsilon \wedge \varepsilon')[l, u]$.

Or. If $KB \models (\phi|\varepsilon)[1, 1]$ and $KB \models (\phi|\varepsilon')[1, 1]$, then $KB \models (\phi|\varepsilon \vee \varepsilon')[1, 1]$.

Another desirable property is *Rational Monotonicity (RM)* (Kraus, Lehmann, & Magidor 1990), which describes a restricted form of monotony, and allows to ignore certain kinds of irrelevant knowledge. The next theorem shows that logical, z -, and *lex-entailment* satisfy *RM*. Here, $KB \not\models C$ denotes that $KB \models C$ does not hold.

Theorem 15 \models , \models^z , and \models^{lex} satisfy the following property for all $KB = (L, P)$ and all events ε , ε' , and ψ :

RM. If $KB \models (\psi|\varepsilon)[1, 1]$ and $KB \not\models (\neg\varepsilon'|\varepsilon)[1, 1]$, then $KB \models (\psi|\varepsilon \wedge \varepsilon')[1, 1]$.

The notion of p -entailment, however, generally does not satisfy *RM*, as the following example shows.

Example 16 Consider the probabilistic knowledge base $KB = (\{bird \Leftarrow eagle\}, \{(fly | bird)[1, 1]\})$. It is easy to see that $(fly | bird)[1, 1]$ is a logical (resp., p -, z -, and *lex*-) consequence of KB , while $(\neg eagle | bird)[1, 1]$ is not a logical (resp., p -, z -, and *lex*-) consequence of KB . Observe now that $(fly | bird \wedge eagle)[1, 1]$ is a logical (resp., z - and *lex*-) consequence of KB , but $(fly | bird \wedge eagle)[1, 1]$ is not a p -consequence of KB . Note that $(fly | bird \wedge eagle)[1, 1]$ is the tight logical (resp., z - and *lex*-) consequence of KB , while $(fly | bird \wedge eagle)[0, 1]$ is the tight p -consequence of KB . \square

We next consider the property *Irrelevance (Irr)* adapted from (Benferhat, Saffiotti, & Smets 2000), which says that ε' is irrelevant to a conclusion " $P \models (\psi|\varepsilon)[1, 1]$ " when they




Figure 1: Relationship between Probabilistic Formalisms and to Classical Formalisms

are defined over disjoint sets of atoms. The following result shows that logical, z -, and lex -entailment satisfy *Irr*.

Theorem 17 \models , \sim^z , and \sim^{lex} satisfy the following property for all $KB = (L, P)$ and all events ε , ε' , and ψ :

Irr. If $KB \sim^s(\psi|\varepsilon)[1, 1]$, and no atom of KB and $(\psi|\varepsilon)[1, 1]$ occurs in ε' , then $KB \sim^s(\psi|\varepsilon\wedge\varepsilon')[1, 1]$.

The notion of p -entailment, however, does not satisfy *Irr*. This is shown by the following example.

Example 18 Consider the probabilistic knowledge base $KB = (\emptyset, \{(fly|bird)[1, 1]\})$. Clearly, $(fly|bird)[1, 1]$ is a logical (resp., p -, z -, and lex -) consequence of KB . Observe now that $(fly|red \wedge bird)[1, 1]$ is a logical (resp., z - and lex -) consequence of KB , but $(fly|red \wedge bird)[1, 1]$ is not a p -consequence of KB . Note that $(fly|red \wedge bird)[1, 1]$ is the tight logical (resp., z - and lex -) consequence of KB , while $(fly|red \wedge bird)[0, 1]$ is the tight p -consequence of KB . \square

We finally consider the property *Direct Inference (DI)* adapted from (Bacchus *et al.* 1996). Informally, *DI* expresses that P should entail all its own conditional constraints (which is similar to *LLE*, but in general not equivalent to *LLE*). The following theorem shows that logical, p -, z -, and lex -entailment all satisfy *DI*.

Theorem 19 \models , \sim^p , \sim^z , and \sim^{lex} satisfy the following property for all probabilistic knowledge bases $KB = (L, P)$, all events ε , ϕ , and ψ , and all $l, u \in [0, 1]$:

DI. If $(\psi|\phi)[l, u] \in P$ and $\varepsilon \Leftrightarrow \phi$ is logically valid, then $KB \sim^s(\psi|\varepsilon)[l, u]$.

Table 3: Summary of Nonmonotonic Properties

Property	\models	\sim^{lex}	\sim^z	\sim^p
KLM postulates	Yes	Yes	Yes	Yes
Rational Monotonicity	Yes	Yes	Yes	No
Irrelevance	Yes	Yes	Yes	No
Direct Inference	Yes	Yes	Yes	Yes

Relationship between Probabilistic Formalisms

We now investigate the relationships between the probabilistic formalisms of this paper. The following theorem shows that logical entailment is stronger than lex -entailment, and that the latter is stronger than z -entailment, which in turn is

stronger than p -entailment. That is, the logical implications illustrated in Fig. 1 hold between the entailment relations.

Theorem 20 Let $KB = (L, P)$ be p -consistent, and let $C = (\psi|\phi)[l, u]$ be a conditional constraint. Then,

(a) $KB \sim^p C$ implies $KB \sim^z C$.

(b) $KB \sim^z C$ implies $KB \sim^{lex} C$.

(c) $KB \sim^{lex} C$ implies $KB \models C$.

In general, none of the converse implications holds, as Table 2 immediately shows. But in the special case where $L \cup P$ has a model in which the conditioning event ϕ has a positive probability, the notions of logical, z -, and lex -entailment of $(\psi|\phi)[l, u]$ from KB all coincide. This important result is expressed by the following theorem.

Theorem 21 Let $KB = (L, P)$ be a p -consistent probabilistic knowledge base, and let $C = (\psi|\phi)[l, u]$ be a conditional constraint such that $L \cup P$ has a model Pr with $Pr(\phi) > 0$. Then, $KB \models C$ iff $KB \sim^{lex} C$ iff $KB \sim^z C$.

The following example shows that p -entailment, however, generally does not coincide with logical entailment when $L \cup P$ has a model Pr with $Pr(\phi) > 0$.

Example 22 Consider again the probabilistic knowledge base $KB_1 = (L_1, P_1)$ shown in Table 1. Then, $L_1 \cup P_1$ has a model Pr with $Pr(eagle) > 0$, and $(legs|eagle)[1, 1]$ is a logical, z -, and lex -consequence of KB , but $(legs|eagle)[1, 1]$ is not a p -consequence of KB . Note that $(legs|eagle)[1, 1]$ is in fact the tight logical, z -, and lex -consequence of KB , while $(legs|eagle)[0, 1]$ is the tight p -consequence of KB . \square

Relationship to Classical Formalisms

We finally explore the relationship between the new notions of p -, z -, and lex -entailment and their classical counterparts. The following theorem shows that the entailment relation \sim^s for p -consistent probabilistic knowledge bases generalizes the classical counterpart \sim^s for p -consistent conditional knowledge bases, where $s \in \{p, z, lex\}$. Here, the operator γ on conditional constraints, sets of conditional constraints, and probabilistic knowledge bases replaces each conditional constraint $(\psi|\phi)[1, 1]$ by the default $\psi \leftarrow \phi$.

Theorem 23 Let $KB = (L, \{(\psi_i|\phi_i)[1, 1] \mid i \in \{1, \dots, n\}\})$ be a p -consistent probabilistic knowledge base, and let $\beta|\alpha$ be a conditional event. Then, for all $s \in \{p, z, lex\}$, it holds that $KB \sim^s(\beta|\alpha)[1, 1]$ iff $\gamma(KB) \sim^s \beta \leftarrow \alpha$.

Algorithms

We now describe algorithms for the main inference problems in weak nonmonotonic probabilistic logics.

Overview

The main decision and optimization problems are as follows:

CONSISTENCY: Given a probabilistic knowledge base KB , decide whether KB is p -consistent.

S-CONSEQUENCE: Given p -consistent probabilistic knowledge base KB and a conditional constraint $(\beta|\alpha)[l, u]$, decide whether $KB \models^s (\beta|\alpha)[l, u]$, for some fixed semantics $s \in \{p, z, lex\}$.

TIGHT S-CONSEQUENCE: Given a p -consistent probabilistic knowledge base KB and a conditional event $\beta|\alpha$, compute $l, u \in [0, 1]$ such that $KB \models^s (\beta|\alpha)[l, u]$, for some fixed semantics $s \in \{p, z, lex\}$.

The basic idea behind the algorithms for solving these decision and optimization problems is to perform a reduction to the following standard decision and optimization problems in model-theoretic probabilistic logic:

POSITIVE PROBABILITY: Given a probabilistic knowledge base $KB = (L, P)$ and an event α , decide whether $L \cup P$ has a model Pr such that $Pr(\alpha) > 0$.

LOGICAL CONSEQUENCE: Given a probabilistic knowledge base KB and a conditional constraint $(\beta|\alpha)[l, u]$, decide whether $KB \models (\beta|\alpha)[l, u]$.

TIGHT LOGICAL CONSEQUENCE: Given a probabilistic knowledge base KB and a conditional event $\beta|\alpha$, compute $l, u \in [0, 1]$ such that $KB \models_{tight} (\beta|\alpha)[l, u]$.

An algorithm for solving the decision problem **CONSISTENCY** (which is similar to the algorithm for deciding ε -consistency in default reasoning by Goldszmidt and Pearl (1991), and which also computes the z -partition of KB , if KB is p -consistent) and an algorithm for solving the optimization problem **TIGHT p -CONSEQUENCE** were presented in (Biazzo *et al.* 2001). The decision problem **p -CONSEQUENCE** can be solved in a similar way.

In the next subsection, we provide algorithms for solving the optimization problems **TIGHT z -** and **TIGHT lex -CONSEQUENCE**. The decision problems **z -** and **lex -CONSEQUENCE** can be solved in a similar way.

Tight S-Consequence

We now present algorithms for solving the optimization problems **TIGHT z -** and **TIGHT lex -CONSEQUENCE**. In the sequel, let $KB = (L, P)$ be a p -consistent probabilistic knowledge base, and let (P_0, \dots, P_k) be its z -partition.

We first provide some preparative definitions as follows. For $G, H \subseteq P$, we say that G is z -preferable to H iff some $i \in \{0, \dots, k\}$ exists such that $P_i \subseteq G$, $P_i \not\subseteq H$, and $P_j \subseteq G$ and $P_j \subseteq H$ for all $i < j \leq k$. We say that G is lex -preferable to H iff some $i \in \{0, \dots, k\}$ exists such that $|G \cap P_i| > |H \cap P_i|$ and $|G \cap P_j| = |H \cap P_j|$ for all $i < j \leq k$. For $\mathcal{D} \subseteq 2^P$ and $s \in \{z, lex\}$, we say G is s -minimal in \mathcal{D} iff $G \in \mathcal{D}$ and no $H \in \mathcal{D}$ is s -preferable to G .

Algorithm tight-z-consequence

Input: p -consistent probabilistic knowledge base $KB = (L, P)$, conditional event $\beta|\alpha$.

Output: interval $[l, u] \subseteq [0, 1]$ such that $KB \models_{tight}^z (\beta|\alpha)[l, u]$.

Notation: (P_0, \dots, P_k) denotes the z -partition of KB .

1. $R := L$;
2. **if** $R \cup \{\alpha > 0\}$ is unsatisfiable **then return** $[1, 0]$;
3. $j := k$;
4. **while** $j \geq 0$ **and** $R \cup P_j \cup \{\alpha > 0\}$ is satisfiable **do begin**
5. $R := R \cup P_j$;
6. $j := j - 1$
7. **end**;
8. compute $l, u \in [0, 1]$ such that $R \models_{tight} (\beta|\alpha)[l, u]$;
9. **return** $[l, u]$.

Figure 2: Algorithm *tight-z-consequence*

Algorithm tight-lex-consequence

Input: p -consistent probabilistic knowledge base $KB = (L, P)$, conditional event $\beta|\alpha$.

Output: interval $[l, u] \subseteq [0, 1]$ such that $KB \models_{tight}^{lex} (\beta|\alpha)[l, u]$.

Notation: (P_0, \dots, P_k) denotes the z -partition of KB .

1. $R := L$;
2. **if** $R \cup \{\alpha > 0\}$ is unsatisfiable **then return** $[1, 0]$;
3. $\mathcal{H} := \{\emptyset\}$;
4. **for** $j := k$ **downto** 0 **do begin**
5. $n := 0$;
6. $\mathcal{H}' := \emptyset$;
7. **for each** $G \subseteq D_j$ **and** $H \in \mathcal{H}$ **do**
8. **if** $R \cup G \cup H \cup \{\alpha > 0\}$ is satisfiable **then**
9. **if** $n = |G|$ **then** $\mathcal{H}' := \mathcal{H}' \cup \{G \cup H\}$
10. **else if** $n < |G|$ **then begin**
11. $\mathcal{H}' := \{G \cup H\}$;
12. $n := |G|$
13. **end**;
14. $\mathcal{H} := \mathcal{H}'$;
15. **end**;
16. $(l, u) := (1, 0)$;
17. **for each** $H \in \mathcal{H}$ **do begin**
18. compute $c, d \in [0, 1]$ s. t. $R \cup H \models_{tight} (\beta|\alpha)[c, d]$;
19. $(l, u) := (\min(l, c), \max(u, d))$
20. **end**;
21. **return** $[l, u]$.

Figure 3: Algorithm *tight-lex-consequence*

The following theorem shows how **TIGHT s -CONSEQUENCE**, where $s \in \{z, lex\}$, can be reduced to **POSITIVE PROBABILITY** and **TIGHT LOGICAL CONSEQUENCE**. The main idea behind this reduction is that there exists a set $\mathcal{D}_\alpha^s(KB) \subseteq 2^P$ such that $KB \models^s (\beta|\alpha)[l, u]$ iff $L \cup H \models (\beta|\alpha)[l, u]$ for all $H \in \mathcal{D}_\alpha^s(KB)$.

Theorem 24 *Let $KB = (L, P)$ be a p -consistent probabilistic knowledge base, and let $\beta|\alpha$ be a conditional event. Let $s \in \{z, lex\}$. Let $\mathcal{D}_\alpha^s(KB)$ denote the set of all s -minimal elements in $\{H \subseteq P \mid L \cup H \cup \{\alpha > 0\} \text{ is satisfiable}\}$. Then, l (resp., u) such that $KB \models_{tight}^s (\beta|\alpha)[l, u]$ is given by:*

- (a) If $L \cup \{\alpha > 0\}$ is unsatisfiable, then $l = 1$ (resp., $u = 0$).
(b) Otherwise, $l = \min c$ (resp., $u = \max d$) subject to $L \cup H \models_{tight} (\beta|\alpha)[c, d]$ and $H \in \mathcal{D}_\alpha^s(KB)$.

For $s = z$ (resp., $s = lex$), Algorithm *tight-s-consequence* (see Fig. 2 (resp., 3)) computes tight intervals under s -entailment. Step 2 checks whether $L \cup \{\alpha > 0\}$ is unsatisfiable. If this is the case, then $[1, 0]$ is returned by Theorem 24 (a). Otherwise, we compute $\mathcal{D}_\alpha^s(KB)$ along the z -partition of KB in steps 3–7 (resp., 3–15), and the requested tight interval using Theorem 24 (b) in step 8 (resp., 16–20).

Computational Complexity

In this section, we draw a precise picture of the computational complexity of the decision and optimization problems described in the previous section.

Complexity Classes

We assume some basic knowledge about the complexity classes P, NP, and co-NP. We now briefly describe some other complexity classes that occur in our results. See especially (Garey & Johnson 1979; Johnson 1990; Papadimitriou 1994) for further background.

The class P^{NP} contains all decision problems that can be solved in deterministic polynomial time with an oracle for NP. The class P_{\parallel}^{NP} contains the decision problems in P^{NP} where all oracle calls must be first prepared and then issued in parallel. The relationship between these complexity classes is described by the following inclusion hierarchy (note that all inclusions are currently believed to be strict):

$$P \subseteq NP, \text{co-NP} \subseteq P_{\parallel}^{NP} \subseteq P^{NP}.$$

To classify problems that compute an output value, rather than a Yes/No-answer, function classes have been introduced. In particular, FP and FP^{NP} are the functional analogs of P and P^{NP} , respectively.

Overview of Complexity Results

We now give an overview of the complexity results. We consider the problems s -CONSEQUENCE and TIGHT s -CONSEQUENCE, where $s \in \{z, lex\}$. Note that CONSISTENCY, p -CONSEQUENCE and TIGHT p -CONSEQUENCE are complete for NP, co-NP, and FP^{NP} , respectively, in the general and in restricted cases (Biazzo *et al.* 2001). We assume that KB and $(\beta|\alpha)[l, u]$ contain only rational numbers.

The complexity results are compactly summarized in Tables 4 and 5. The problems z - and lex -CONSEQUENCE are complete for the classes P_{\parallel}^{NP} and P^{NP} , respectively, whereas the problems TIGHT z - and TIGHT lex -CONSEQUENCE are both complete for the class FP^{NP} .

The hardness often holds even in the restricted *literal-Horn case*, where KB and $\beta|\alpha$ are both literal-Horn. Here, a conditional event $\psi|\phi$ (resp., logical constraint $\psi \Leftarrow \phi$) is literal-Horn iff ψ is a basic event (resp., ψ is either a basic event or the negation of a basic event) and ϕ is either \top or a conjunction of basic events. A conditional constraint $(\psi|\phi)[l, u]$ is literal-Horn iff the conditional event $\psi|\phi$ is

literal-Horn. A probabilistic knowledge base $KB = (L, P)$ is literal-Horn iff each member of $L \cup P$ is literal-Horn.

Table 4: Complexity of s -CONSEQUENCE

Problem	Complexity
z -CONSEQUENCE	P_{\parallel}^{NP} -complete
lex -CONSEQUENCE	P^{NP} -complete

Table 5: Complexity of TIGHT s -CONSEQUENCE

Problem	Complexity
TIGHT z -CONSEQUENCE	FP^{NP} -complete
TIGHT lex -CONSEQUENCE	FP^{NP} -complete

Detailed Complexity Results

The following theorem shows that z - and lex -CONSEQUENCE are complete for the classes P_{\parallel}^{NP} and P^{NP} , respectively. Here, hardness for P_{\parallel}^{NP} and P^{NP} follows from Theorem 23 and the P_{\parallel}^{NP} - and P^{NP} -hardness of deciding Pearl's entailment in System Z and Lehmann's lexicographic entailment (Eiter & Lukasiewicz 2000).

Theorem 25 *Given a p -consistent probabilistic knowledge base $KB = (L, P)$, and a conditional constraint $(\beta|\alpha)[l, u]$, deciding whether $KB \models^z (\beta|\alpha)[l, u]$, where $s = z$ (resp., $s = lex$) is complete for P_{\parallel}^{NP} (resp., P^{NP}). For $s = lex$, hardness holds even if KB and $\beta|\alpha$ are both literal-Horn.*

The next theorem shows that TIGHT s -CONSEQUENCE, where $s \in \{z, lex\}$, is FP^{NP} -complete. Here, hardness holds by a polynomial reduction from the FP^{NP} -complete *traveling salesman cost* problem (Papadimitriou 1994).

Theorem 26 *Given a p -consistent probabilistic knowledge base $KB = (L, P)$, and a conditional event $\beta|\alpha$, computing $l, u \in [0, 1]$ such that $KB \models_{tight}^s (\beta|\alpha)[l, u]$, where $s \in \{z, lex\}$, is complete for FP^{NP} . Hardness holds even if KB and $\beta|\alpha$ are both literal-Horn, and $L = \emptyset$.*

Related Work

We now describe the relationship to probabilistic logic under coherence and to strong nonmonotonic probabilistic logics.

Probabilistic Logic under Coherence

The notions of p -consistency and p -entailment coincide with the notions of g -coherence and g -coherent entailment, respectively, from probabilistic logic under coherence.

Probabilistic reasoning under coherence is an approach to reasoning with conditional constraints, which has been extensively explored especially in the field of statistics, and which is based on the coherence principle of de Finetti and suitable generalizations of it (see, for example, the work by Biazzo and Gilio (2000), Gilio (1995; 2002), and Gilio and Scozzafava (1994)), or on similar principles that have been

adopted for lower and upper probabilities (Pelessoni and Vici (1998), Vici (1996), and Walley (1991)). We now recall the main concepts from probabilistic logic under coherence, and then formulate the above equivalence results.

We first define (precise) probability assessments and their coherence. A *probability assessment* (L, A) on a set of conditional events \mathcal{E} consists of a set of logical constraints L , and a mapping A from \mathcal{E} to $[0, 1]$. Informally, L describes logical relationships, while A represents probabilistic knowledge. For $\{\psi_1|\phi_1, \dots, \psi_n|\phi_n\} \subseteq \mathcal{E}$ with $n \geq 1$ and n real numbers s_1, \dots, s_n , let the mapping $G: \mathcal{I}_{\Phi} \rightarrow \mathbf{R}$ be defined as follows. For every $I \in \mathcal{I}_{\Phi}$:

$$G(I) = \sum_{i=1}^n s_i \cdot I(\phi_i) \cdot (I(\psi_i) - A(\psi_i|\phi_i)).$$

In the framework of betting criterion, G can be interpreted as the random gain corresponding to a combination of n bets of amounts $s_1 \cdot A(\psi_1|\phi_1), \dots, s_n \cdot A(\psi_n|\phi_n)$ on $\psi_1|\phi_1, \dots, \psi_n|\phi_n$ with stakes s_1, \dots, s_n . More precisely, to bet on $\psi_i|\phi_i$, one pays an amount of $s_i \cdot A(\psi_i|\phi_i)$, and one gets back the amounts of $s_i, 0$, and $s_i \cdot A(\psi_i|\phi_i)$, when $\psi_i \wedge \phi_i, \neg\psi_i \wedge \phi_i$, and $\neg\phi_i$, respectively, turn out to be true. The following notion of *coherence* assures that it is impossible (for both the gambler and the bookmaker) to have sure (or uniform) loss. A probability assessment (L, A) on a set of conditional events \mathcal{E} is *coherent* iff for every $\{\psi_1|\phi_1, \dots, \psi_n|\phi_n\} \subseteq \mathcal{E}, n \geq 1$, and for all reals s_1, \dots, s_n , it holds $\max\{G(I) \mid I \in \mathcal{I}_{\Phi}, I \models L \cup \{\phi_1 \vee \dots \vee \phi_n\}\} \geq 0$.

We next define imprecise probability assessments and the notions of g-coherence and of g-coherent entailment for them. An *imprecise probability assessment* (L, A) on a set of conditional events \mathcal{E} consists of a set of logical constraints L and a mapping A that assigns to each $\varepsilon \in \mathcal{E}$ an interval $[l, u] \subseteq [0, 1], l \leq u$. We say (L, A) is *g-coherent* iff a coherent precise probability assessment (L, A^*) on \mathcal{E} exists with $A^*(\varepsilon) \in A(\varepsilon)$ for all $\varepsilon \in \mathcal{E}$. The imprecise probability assessment $[l, u]$ on a conditional event γ , denoted $\{(\gamma, [l, u])\}$, is called a *g-coherent consequence* of (L, A) iff $A^*(\gamma) \in [l, u]$ for every g-coherent precise probability assessment A^* on $\mathcal{E} \cup \{\gamma\}$ such that $A^*(\varepsilon) \in A(\varepsilon)$ for all $\varepsilon \in \mathcal{E}$. It is a *tight g-coherent consequence* of (L, A) iff l (resp., u) is the infimum (resp., supremum) of $A^*(\gamma)$ subject to all g-coherent precise probability assessments A^* on $\mathcal{E} \cup \{\gamma\}$ such that $A^*(\varepsilon) \in A(\varepsilon)$ for all $\varepsilon \in \mathcal{E}$.

We finally define the concepts of g-coherence and of g-coherent entailment for probabilistic knowledge bases (Biazzo *et al.* 2002). Every imprecise probability assessment $IP = (L, A)$, where L is finite, and A is defined on a finite set of conditional events \mathcal{E} , can be represented by a probabilistic knowledge base. Conversely, every *reduced* probabilistic knowledge base $KB = (L, P)$, where (i) $l \leq u$ for all $(\varepsilon)[l, u] \in P$, and (ii) $\varepsilon_1 \neq \varepsilon_2$ for any two distinct $(\varepsilon_1)[l_1, u_1], (\varepsilon_2)[l_2, u_2] \in P$, can be expressed by the imprecise assessment $IP_{KB} = (L, A_{KB})$ on \mathcal{E}_{KB} , where

$$A_{KB} = \{(\psi|\phi, [l, u]) \mid (\psi|\phi)[l, u] \in KB\},$$

$$\mathcal{E}_{KB} = \{\psi|\phi \mid \exists l, u \in [0, 1]: (\psi|\phi)[l, u] \in KB\}.$$

A reduced probabilistic knowledge base KB is *g-coherent* iff IP_{KB} is g-coherent. In this case, a conditional constraint $(\psi|\phi)[l, u]$ is a *g-coherent* (resp., *tight g-coherent*)

consequence of KB , denoted $KB \Vdash^g(\psi|\phi)[l, u]$ (resp., $KB \Vdash_{tight}^g(\psi|\phi)[l, u]$), iff $\{(\psi|\phi, [l, u])\}$ is a g-coherent (resp., tight g-coherent) consequence of IP_{KB} .

The following theorem shows that g-coherence and g-coherent entailment coincide with p -consistency and p -entailment, respectively. It follows immediately from Theorems 5 and 7 and similar characterizations of g-coherence and g-coherent entailment through conditional constraint rankings due to Biazzo *et al.* (2002).

Theorem 27 *Let $KB = (L, P)$ be a reduced probabilistic knowledge base, and let C be a conditional constraint. Then, (a) KB is g-coherent iff KB is p -consistent; and (b) if KB is p -consistent, then $KB \Vdash^g C$ iff $KB \Vdash^p C$.*

Strong Nonmonotonic Probabilistic Logics

A companion paper (Lukasiewicz 2002) presents similar probabilistic generalizations of Pearl's entailment in System Z and of Lehmann's lexicographic entailment, which are, however, quite different from the ones in this paper.

More precisely, the formalisms presented in (Lukasiewicz 2002) add to logical entailment in model-theoretic probabilistic logic (i) some inheritance of purely probabilistic knowledge, and (ii) a strategy for resolving inconsistencies due to the inheritance of logical and purely probabilistic knowledge. For this reason, they are generally much stronger than logical entailment. Thus, they are especially useful where logical entailment is too weak, for example, in probabilistic logic programming (Lukasiewicz 2001b; 2001a) and probabilistic ontology reasoning in the Semantic Web (Giugno & Lukasiewicz 2002). Other applications are deriving degrees of belief from statistical knowledge and degrees of belief, handling inconsistencies in probabilistic knowledge bases, and probabilistic belief revision.

In particular, in reasoning from statistical knowledge and degrees of belief, they show a similar behavior as reference-class reasoning (Reichenbach 1949; Kyburg, Jr. 1974; 1983; Pollock 1990) in a number of uncontroversial examples. However, they also avoid many drawbacks of reference-class reasoning (Lukasiewicz 2002): They can handle complex scenarios and even purely probabilistic subjective knowledge as input. Furthermore, conclusions are drawn in a global way from all the available knowledge as a whole. The following example illustrates the use of *strong lex-entailment* (Lukasiewicz 2002) for reasoning from statistical knowledge and degrees of belief.

Example 28 Suppose that we have the statistical knowledge “all penguins are birds”, “between 90% and 95% of all birds fly”, “at most 5% of all penguins fly”, and “at least 95% of all yellow objects are easy to see”. Moreover, assume that we believe “Sam is a yellow penguin”. What do we then conclude about Sam's property of being easy to see? Under reference-class reasoning, which is a machinery for dealing with such statistical knowledge and degrees of belief, we conclude “Sam is easy to see with a probability of at least 0.95”. This is also what we obtain using the notion of strong *lex-entailment*: The above statistical knowledge can be represented by the probabilistic knowledge base $KB = (L, P) = (\{bird \Leftarrow penguin\}, \{(fly \mid bird)[0.9, 0.95]$,

$(fly \mid penguin)[0, 0.05], (easy_to_see \mid yellow)[0.95, 1]$). It is then easy to verify that KB is *strongly p-consistent*, and that under strong *lex-entailment* from KB , we obtain the tight conclusion $(easy_to_see \mid yellow \wedge penguin)[0.95, 1]$, as desired; see (Lukasiewicz 2002).

Notice that KB is also satisfiable and *p-consistent*, and under logical and *p*-, *z*-, and *lex-entailment* from KB , we have $(easy_to_see \mid yellow \wedge penguin)[0, 1]$, rather than the above conditional constraint, as tight conclusion. \square

Summary and Conclusion

Towards probabilistic formalisms for resolving local inconsistencies under model-theoretic probabilistic entailment, we have introduced novel probabilistic generalizations of Pearl's entailment in System Z and of Lehmann's lexicographic entailment. We have then analyzed the nonmonotonic and semantic properties of the new notions of probabilistic entailment. Furthermore, we have presented algorithms for reasoning under the new formalisms, and we have given a precise picture of its computational complexity.

As an important feature of the new notions of entailment in System Z and of lexicographic entailment, we have shown that they coincide with model-theoretic probabilistic entailment whenever there are no local inconsistencies. That is, the new formalisms are essentially identical to model-theoretic probabilistic entailment, except that they resolve the problem of local inconsistencies. In particular, this property also distinguishes the new notions of entailment in this paper from the notion of probabilistic entailment under coherence and from the notions of entailment in strong nonmonotonic probabilistic logics (Lukasiewicz 2002).

More precisely, probabilistic entailment under coherence is related to the new formalisms in this paper, since it is a generalization of default reasoning in System P (see also (Biazzo *et al.* 2002; 2001)). However, there are several crucial differences. First, the formalisms in this paper are generalizations of the more sophisticated notions of entailment in System Z and lexicographic entailment, rather than entailment in System P . As a consequence, they have nicer semantic properties, and are strictly stronger than probabilistic entailment under coherence. Second, as for resolving local inconsistencies as described in Example 1, the formalisms here coincide with model-theoretic probabilistic entailment whenever there are no local inconsistencies, while probabilistic entailment under coherence does not.

The notions of entailment in strong nonmonotonic probabilistic logics (Lukasiewicz 2002), in contrast, aim at increasing the inferential power of model-theoretic probabilistic entailment by adding some restricted forms of *P-INH* (recall that model-theoretic probabilistic entailment completely lacks *P-INH*). For this reason, the notions of entailment in (Lukasiewicz 2002) are generally much stronger than model-theoretic probabilistic entailment. For example, under the notions of entailment in (Lukasiewicz 2002), we can conclude $(fly \mid eagle)[0.95, 1]$ from KB_1 of Table 1.

An interesting topic of future research is to develop and explore further nonmonotonic formalisms for reasoning with conditional constraints. Besides extending classical formalisms for default reasoning, which may addition-

ally contain a strength assignment to the defaults, one may also think about combining the new formalisms of this paper and of (Lukasiewicz 2002) with some probability selection technique (e.g., maximum entropy or center of mass).

Acknowledgments. This work has been supported by the Austrian Science Fund project Z29-N04 and a Marie Curie Individual Fellowship of the European Community programme "Human Potential" under contract number HPMF-CT-2001-001286 (disclaimer: The author is solely responsible for information communicated and the European Commission is not responsible for any views or results expressed). I am grateful to Angelo Gilio and Lluís Godó for useful comments on an earlier version of this paper.

References

- Adams, E. W. 1975. *The Logic of Conditionals*, volume 86 of *Synthese Library*. Dordrecht, Netherlands: D. Reidel.
- Amarger, S.; Dubois, D.; and Prade, H. 1991. Constraint propagation with imprecise conditional probabilities. In *Proceedings UAI-91*, 26–34. Morgan Kaufmann.
- Bacchus, F.; Grove, A.; Halpern, J. Y.; and Koller, D. 1996. From statistical knowledge bases to degrees of belief. *Artif. Intell.* 87(1–2):75–143.
- Benferhat, S.; Cayrol, C.; Dubois, D.; Lang, J.; and Prade, H. 1993. Inconsistency management and prioritized syntax-based entailment. In *Proceedings IJCAI-93*, 640–645. Morgan Kaufmann.
- Benferhat, S.; Dubois, D.; and Prade, H. 1992. Representing default rules in possibilistic logic. In *Proceedings KR-92*, 673–684. Morgan Kaufmann.
- Benferhat, S.; Dubois, D.; and Prade, H. 1997. Nonmonotonic reasoning, conditional objects and possibility theory. *Artif. Intell.* 92(1–2):259–276.
- Benferhat, S.; Saffiotti, A.; and Smets, P. 2000. Belief functions and default reasoning. *Artif. Intell.* 122(1–2):1–69.
- Biazzo, V., and Gilio, A. 2000. A generalization of the fundamental theorem of de Finetti for imprecise conditional probability assessments. *Int. J. Approx. Reasoning* 24(2–3):251–272.
- Biazzo, V.; Gilio, A.; Lukasiewicz, T.; and Sanfilippo, G. 2001. Probabilistic logic under coherence: Complexity and algorithms. In *Proceedings ISIPTA-01*, 51–61. Extended Report INFSYS RR-1843-01-04, Institut für Informationssysteme, Technische Universität Wien, 2001.
- Biazzo, V.; Gilio, A.; Lukasiewicz, T.; and Sanfilippo, G. 2002. Probabilistic logic under coherence, model-theoretic probabilistic logic, and default reasoning in System P . *Journal of Applied Non-Classical Logics* 12(2):189–213.
- Boole, G. 1854. *An Investigation of the Laws of Thought, on which are Founded the Mathematical Theories of Logic and Probabilities*. London: Walton and Maberley. Reprint: Dover Publications, New York, 1958.
- Dubois, D., and Prade, H. 1988. On fuzzy syllogisms. *Computational Intelligence* 4(2):171–179.

- Dubois, D., and Prade, H. 1991. Possibilistic logic, preferential models, non-monotonicity and related issues. In *Proceedings IJCAI-91*, 419–424. Morgan Kaufmann.
- Dubois, D., and Prade, H. 1994. Conditional objects as nonmonotonic consequence relationships. *IEEE Trans. Syst. Man Cybern.* 24(12):1724–1740.
- Eiter, T., and Lukasiewicz, T. 2000. Default reasoning from conditional knowledge bases: Complexity and tractable cases. *Artif. Intell.* 124(2):169–241.
- Fagin, R.; Halpern, J. Y.; and Megiddo, N. 1990. A logic for reasoning about probabilities. *Inf. Comput.* 87:78–128.
- Frisch, A. M., and Haddawy, P. 1994. Anytime deduction for probabilistic logic. *Artif. Intell.* 69(1–2):93–122.
- Gabbay, D. M., and Smets, P., eds. 1998. *Handbook on Defeasible Reasoning and Uncertainty Management Systems*. Dordrecht, Netherlands: Kluwer Academic.
- Garey, M. R., and Johnson, D. S. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York: Freeman.
- Geffner, H. 1992. *Default Reasoning: Causal and Conditional Theories*. MIT Press.
- Gilio, A., and Scozzafava, R. 1994. Conditional events in probability assessment and revision. *IEEE Trans. Syst. Man Cybern.* 24(12):1741–1746.
- Gilio, A. 1995. Probabilistic consistency of conditional probability bounds. In *Advances in Intelligent Computing*, volume 945 of *LNCS*, 200–209. Springer.
- Gilio, A. 2002. Probabilistic reasoning under coherence in System *P*. *Ann. Math. Artif. Intell.* 34(1–3):5–34.
- Giugno, R., and Lukasiewicz, T. 2002. $P\text{-}SHOQ(\mathbf{D})$: A probabilistic extension of $SHOQ(\mathbf{D})$ for probabilistic ontologies in the Semantic Web. In *Proceedings JELIA-02*, volume 2424 of *LNCS/LNAI*, 86–97. Springer.
- Goldszmidt, M., and Pearl, J. 1991. On the consistency of defeasible databases. *Artif. Intell.* 52(2):121–149.
- Goldszmidt, M., and Pearl, J. 1992. Rank-based systems: A simple approach to belief revision, belief update and reasoning about evidence and actions. In *Proceedings KR-92*, 661–672. Morgan Kaufmann.
- Goldszmidt, M., and Pearl, J. 1996. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artif. Intell.* 84(1–2):57–112.
- Hailperin, T. 1996. *Sentential Probability Logic: Origins, Development, Current Status, and Technical Applications*. London, UK: Associated University Presses.
- Johnson, D. S. 1990. A catalog of complexity classes. In van Leeuwen, J., ed., *Handbook of Theoretical Computer Science*, volume A. MIT Press. chapter 2, 67–161.
- Kraus, S.; Lehmann, D.; and Magidor, M. 1990. Non-monotonic reasoning, preferential models and cumulative logics. *Artif. Intell.* 14(1):167–207.
- Kyburg, Jr., H. E. 1974. *The Logical Foundations of Statistical Inference*. Dordrecht, Netherlands: D. Reidel.
- Kyburg, Jr., H. E. 1983. The reference class. *Philos. Sci.* 50:374–397.
- Lamarre, P. 1992. A promenade from monotonicity to non-monotonicity following a theorem prover. In *Proceedings KR-92*, 572–580. Morgan Kaufmann.
- Lehmann, D., and Magidor, M. 1992. What does a conditional knowledge base entail? *Artif. Intell.* 55(1):1–60.
- Lehmann, D. 1989. What does a conditional knowledge base entail? In *Proc. KR-89*, 212–222. Morgan Kaufmann.
- Lehmann, D. 1995. Another perspective on default reasoning. *Ann. Math. Artif. Intell.* 15(1):61–82.
- Lukasiewicz, T. 1999a. Local probabilistic deduction from taxonomic and probabilistic knowledge-bases over conjunctive events. *Int. J. Approx. Reasoning* 21(1):23–61.
- Lukasiewicz, T. 1999b. Probabilistic deduction with conditional constraints over basic events. *J. Artif. Intell. Res.* 10:199–241.
- Lukasiewicz, T. 2001a. Probabilistic logic programming under inheritance with overriding. In *Proceedings UAI-01*, 329–336. Morgan Kaufmann.
- Lukasiewicz, T. 2001b. Probabilistic logic programming with conditional constraints. *ACM Transactions on Computational Logic (TOCL)* 2(3):289–339.
- Lukasiewicz, T. 2002. Probabilistic default reasoning with conditional constraints. *Ann. Math. Artif. Intell.* 34(1–3):35–88.
- Lukasiewicz, T. 2003. Weak nonmonotonic probabilistic logics. Technical Report INFSYS RR-1843-02-02, Institut für Informationssysteme, Technische Universität Wien.
- Nilsson, N. J. 1986. Probabilistic logic. *Artif. Intell.* 28(1):71–88.
- Papadimitriou, C. H. 1994. *Computational Complexity*. Reading, MA: Addison-Wesley.
- Pearl, J. 1989. Probabilistic semantics for nonmonotonic reasoning: A survey. In *Proceedings KR-89*, 505–516. Morgan Kaufmann.
- Pearl, J. 1990. System *Z*: A natural ordering of defaults with tractable applications to default reasoning. In *Proceedings TARK-90*, 121–135. Morgan Kaufmann.
- Pelessoni, R., and Vicig, P. 1998. A consistency problem for imprecise conditional probability assessments. In *Proceedings IPMU-98*, 1478–1485.
- Pollock, J. L. 1990. *Nomic Probabilities and the Foundations of Induction*. Oxford: Oxford University Press.
- Reichenbach, H. 1949. *Theory of Probability*. Berkeley, CA: University of California Press.
- Shoham, Y. 1987. A semantical approach to nonmonotonic logics. In *Proceedings LICS-87*, 275–279.
- Spohn, W. 1988. Ordinal conditional functions: A dynamic theory of epistemic states. In Harper, W., and Skyrms, B., eds., *Causation in Decision, Belief Change, and Statistics*, volume 2. Dordrecht, Netherlands: Reidel. 105–134.
- Vicig, P. 1996. An algorithm for imprecise conditional probability assessments in expert systems. In *Proceedings IPMU-96*, 61–66.
- Walley, P. 1991. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall.