

Probabilistic Reasoning with Inconsistent Beliefs Using Inconsistency Measures

Nico Potyka

FernUniversität in Hagen, Germany
nico.potyka@FernUni-Hagen.de

Matthias Thimm

University of Koblenz-Landau, Germany
thimm@uni-koblenz.de

Abstract

The classical probabilistic entailment problem is to determine upper and lower bounds on the probability of formulas, given a consistent set of probabilistic assertions. We generalize this problem by omitting the consistency assumption and, thus, provide a general framework for probabilistic reasoning under inconsistency. To do so, we utilize inconsistency measures to determine probability functions that are closest to satisfying the knowledge base. We illustrate our approach on several examples and show that it has both nice formal and computational properties.

1 Introduction

Many branches in artificial intelligence deal with reasoning under uncertainty and inconsistency, e. g., default reasoning [Reiter, 1980], paraconsistent logics [Béziau *et al.*, 2007], belief dynamics [Hansson, 2001], computational argumentation [Bench-Capon and Dunne, 2007] and probabilistic reasoning [Nilsson, 1986]. Inconsistencies arise easily in many applications, e. g., when several experts share their knowledge in order to solve a problem [Konieczny and Perez, 2011].

We consider the scenario that our knowledge is both *uncertain* and *inconsistent*. As a simple example, consider two experts, the one arguing that the price of a stock will probably rise, the other arguing that the price will probably fall. Even though uncertain, taken together both statements are inconsistent. How can a rational agent incorporate both beliefs simultaneously?

To represent uncertain knowledge, we use an extension of classical probabilistic logic [Nilsson, 1986] and consider probabilistic conditionals $(\psi | \phi)[d]$ that encode uncertain rules 'if ϕ then ψ with probability d ' [Benferhat *et al.*, 1999; Kern-Isberner, 2001]. Inconsistencies occur in this framework when multiple conditionals cannot be satisfied jointly by a probability function. To deal with inconsistencies, we generalize the probabilistic entailment problem [Jaumard *et al.*, 1991; Lukasiewicz, 1999] to inconsistent knowledge bases by using inconsistency measures [Grant and Hunter, 2013b]. An inconsistency measure \mathcal{I} is a function that maps a knowledge base to a non-negative real number such that larger values indicate larger inconsistency. For probabilistic

conditional logic, several inconsistency measures have been proposed, see, e. g. [Thimm, 2013; Picado-Muiño, 2011]. We apply the family of minimal violation measures from [Potyka, 2014] since they allow us to extend the classical notion of models of a probabilistic knowledge base to inconsistent ones. Intuitively, the generalized models are those probability functions that *minimally violate* the knowledge base [Potyka and Thimm, 2014]. We incorporate integrity constraints and study a family of generalized entailment problems for probabilistic knowledge bases. More specifically, the contributions of this work are as follows:

1. We introduce the computational problem of generalized entailment with integrity constraints in probabilistic logics and thus provide an approach to reasoning with inconsistent probabilistic knowledge (Section 3).
2. We analyse the behaviour of our approach by showing that it satisfies several rationality postulates (Section 4).
3. We show how to solve the generalized entailment problem and that this is computationally not harder than solving the classical probabilistic entailment problem for consistent knowledge bases (Section 5).

We explain the necessary basics in Section 2, discuss related work in Section 6, and conclude in Section 7.

2 Preliminaries

We consider a propositional language $\mathcal{L}(\text{At})$ built up over a finite set of propositional variables At in the usual way. For $\phi, \psi \in \mathcal{L}(\text{At})$ we abbreviate $\phi \wedge \psi$ by $\phi\psi$ and $\neg\phi$ by $\bar{\phi}$.

A *possible world* assigns a truth value to each $a \in \text{At}$. Let $\Omega(\text{At})$ denote the set of all possible worlds. $\omega \in \Omega(\text{At})$ *satisfies* an atom $a \in \text{At}$, denoted by $\omega \models a$, if and only if $\omega(a) = \text{true}$. \models is extended to complex formulas in $\mathcal{L}(\text{At})$ in the usual way. Formulas $\psi, \phi \in \mathcal{L}(\text{At})$ are *equivalent*, denoted by $\phi \equiv \psi$, if and only if $\omega \models \phi$ whenever $\omega \models \psi$ for every $\omega \in \Omega(\text{At})$ and vice versa.

We build up a probabilistic language $(\mathcal{L}(\text{At}) | \mathcal{L}(\text{At}))^{pr}$ containing a *probabilistic conditional* $(\psi | \phi)[d]$ for all $\phi, \psi \in \mathcal{L}(\text{At})$ and $d \in [0, 1]$. Intuitively, $(\psi | \phi)[d]$ says that if ϕ is true then ψ is also true with probability d (see below). If ϕ is tautological, $\phi \equiv \top$, we abbreviate $(\psi | \phi)[d]$ by $(\psi)[d]$.

A *knowledge base* \mathcal{K} is an ordered finite subset of $(\mathcal{L}(\text{At}) | \mathcal{L}(\text{At}))^{pr}$. We impose an ordering on the conditionals in a knowledge base only for technical convenience. The

order can be arbitrary and has no further meaning other than to enumerate the conditionals of a knowledge base in an unambiguous way.

Semantics are given to probabilistic conditionals by *probability functions* over $\Omega(\text{At})$, which are denoted by $\mathcal{P}(\text{At})$. The probability of a formula $\phi \in \mathcal{L}(\text{At})$ with respect to $P \in \mathcal{P}(\text{At})$ is defined by $P(\phi) = \sum_{\omega \models \phi} P(\omega)$. As usual in this context, P satisfies a probabilistic conditional $(\psi | \phi)[d]$, denoted by $P \models^{pr} (\psi | \phi)[d]$, if and only if $P(\psi\phi) = dP(\phi)$ [Nilsson, 1986; Paris, 1994]. A probability function P satisfies a knowledge base \mathcal{K} (or is a model of \mathcal{K}), denoted by $P \models^{pr} \mathcal{K}$, if and only if $P \models^{pr} c$ for every $c \in \mathcal{K}$. Let $\text{Mod}(\mathcal{K}) \subseteq \mathcal{P}(\text{At})$ be the set of models of \mathcal{K} . If $\text{Mod}(\mathcal{K}) = \emptyset$ then \mathcal{K} is called *inconsistent*.

Broadly speaking, there are two main approaches to reason with probabilistic logics. First, we can consider the whole set of models $\text{Mod}(\mathcal{K})$ of \mathcal{K} and use it to derive probability intervals for given formulas [Nilsson, 1986; Jaumard *et al.*, 1991]. Second, we can search for a best model $P^* \in \text{Mod}(\mathcal{K})$ with respect to some common sense rationales and use P^* to compute the probabilities directly [Nilsson, 1986; Paris, 1994; Kern-Isberner, 2001]. However, if \mathcal{K} is inconsistent, there is no way to infer reasonable information with these approaches because there exists no model at all.

Inconsistency measures help analyzing inconsistent knowledge bases by assigning nonnegative values to knowledge bases that quantify the degree of inconsistency, see, e. g., [Knight, 2002; Hunter and Konieczny, 2010; Thimm, 2013]. The family of *minimal violation measures* is defined by measuring the violation of the equations defined by the probabilistic satisfaction relation [Potyka, 2014]. To understand how, note that the condition $P(\psi\phi) = dP(\phi)$ is a linear constraint over P . With a slight abuse of notation, let us identify P with a probability vector $(P(\omega_1) \dots P(\omega_n))$, $n = |\Omega(\text{At})|$; and for a formula F , let the indicator function $1_{\{F\}}(\omega)$ map to 1 iff $\omega \models F$ and to 0 otherwise. Then we can rewrite $P(\psi\phi) = dP(\phi)$ in vector notation as $a_c P = 0$, where a_c is the transpose of the vector

$$(1_{\{\psi\phi\}}(\omega_j) \cdot (1-d) - 1_{\{\bar{\psi}\phi\}}(\omega_j) \cdot d)_{1 \leq j \leq n},$$

see also [Nilsson, 1986; Jaumard *et al.*, 1991]. Now given a knowledge base \mathcal{K} , we associate \mathcal{K} with the $(m \times n)$ -matrix

$$A_{\mathcal{K}} = \begin{pmatrix} a_1 \\ \dots \\ a_m \end{pmatrix}.$$

The linear equation system $A_{\mathcal{K}}x = 0$ can be solved by a probability vector P if and only if \mathcal{K} is consistent, see also [Nilsson, 1986; Jaumard *et al.*, 1991]. The minimal violation value $\mathcal{I}_{\Pi}^p(\mathcal{K})$ of \mathcal{K} with respect to the minimal violation measure \mathcal{I}_{Π}^p is the solution of the following optimization problem

$$\begin{aligned} & \min_{x \in \mathbb{R}^n} && \|A_{\mathcal{K}}x\|_p && (1) \\ & \text{subject to} && \sum_{i=1}^n x_i = 1 \\ & && x \geq 0, \end{aligned}$$

where $\|\cdot\|_p$ denotes the p -norm defined by $\|x\|_p = \sqrt[p]{\sum_{i=1}^n |x_i|^p}$ for $p \geq 1$. Well-known special cases are the *1-norm* $\|x\|_1 = \sum_{i=1}^n |x_i|$, the *Euclidean norm* $\|x\|_2 = \sqrt{\sum_{i=1}^n x_i^2}$ and the limit for $p \rightarrow \infty$, the *maximum norm* $\|x\|_{\infty} = \max\{|x_i| \mid 1 \leq i \leq n\}$.

Note that the constraints of (1) guarantee that each feasible solution is a probability vector. By definiteness of norms, it holds $\mathcal{I}_{\Pi}^p(\mathcal{K}) = 0$ iff there is a P such that $A_{\mathcal{K}}P = 0$, i. e., iff \mathcal{K} is consistent. As \mathcal{K} becomes 'more' inconsistent, $\mathcal{I}_{\Pi}^p(\mathcal{K})$ increases continuously, see [Potyka, 2014] for the details and further properties.

The probability functions minimizing (1) can be regarded to be as close as possible to a model of \mathcal{K} in the sense that they minimally violate the corresponding equation system. In fact, if \mathcal{K} is consistent, they correspond to the models of \mathcal{K} and are therefore called *generalized models* of \mathcal{K} [Potyka and Thimm, 2014]. More formally, the set of generalized models is defined as follows:

$$\text{GMod}^p(\mathcal{K}) = \{P \in \mathcal{P}(\text{At}) \mid \|A_{\mathcal{K}}P\|_p = \mathcal{I}_{\Pi}^p(\mathcal{K})\}.$$

In [Potyka and Thimm, 2014] generalized maximum entropy reasoning is considered. That is, among all generalized models one selects the one maximizing entropy. The generalized maximum entropy model can be used to repair the knowledge base or to compute probabilities for arbitrary formulas. This approach has some nice properties and can be computed by convex programming techniques [Potyka and Thimm, 2014].

3 Generalized Entailment with Integrity Constraints

We will now focus on generalizing the second major approach to reason with consistent knowledge, namely reasoning with all models. This problem is usually called the *probabilistic entailment problem* [Jaumard *et al.*, 1991].

Given a consistent knowledge base \mathcal{K} and a query $(\psi | \phi)$, $\phi, \psi \in \mathcal{L}(\text{At})$, the probabilistic entailment problem is to find a *tight* probability interval $[l, u]$ such that $P(\psi | \phi) \in [l, u]$ for all $P \in \text{Mod}(\mathcal{K})$ with $P(\phi) > 0$ [Jaumard *et al.*, 1991; Lukasiewicz, 1999]. 'Tight' means that the probability interval cannot be further decreased without violating the condition [Lukasiewicz, 1999]. This condition is important for otherwise the interval $[0, 1]$ always yields a feasible and completely non-informative solution. We denote the classical probabilistic entailment relation by \models^c , i. e., if $[l, u]$ is the corresponding tight probability interval, we write $\mathcal{K} \models^c (\psi | \phi)[l, u]$. If there is no $P \in \text{Mod}(\mathcal{K})$ with $P(\phi) > 0$, we follow [Lukasiewicz, 1999] and let $l = 1, u = 0$.

By replacing $\text{Mod}(\mathcal{K})$, with the generalized models $\text{GMod}^p(\mathcal{K})$, the *generalized entailment problem* can be defined. The lower and upper bounds l and u can be obtained by solving the two optimization problems

$$\begin{aligned} & \text{opt}_{P \in \text{GMod}^p(\mathcal{K})} && P(\psi | \phi) && (2) \\ & \text{subject to} && P(\phi) > 0, \end{aligned}$$

where opt stands for min and max, respectively. We want to consider a slightly more general problem. In addition to our knowledge base \mathcal{K} , which might be inconsistent, we consider

a second knowledge base \mathcal{IC} which is assumed to be consistent. The conditionals in \mathcal{IC} are called *integrity constraints*. To begin with, we generalize some basic concepts.

Definition 1 (Minimal Violation Measures with Integrity Constraints). The minimal violation value $\mathcal{I}_{\mathcal{IC}}^p(\mathcal{K})$ of \mathcal{K} with respect to the minimal violation measure $\mathcal{I}_{\mathcal{IC}}^p$ with integrity constraints \mathcal{IC} is the solution of the optimization problem

$$\min_{P \in \text{Mod}(\mathcal{IC})} \|A_{\mathcal{K}}P\|_p \quad (3)$$

Definition 2 (Generalized Models with Integrity Constraints). The set of probability functions minimizing (3) is called the set of generalized models of \mathcal{K} with respect to the integrity constraints \mathcal{IC} and is denoted by $\text{GMod}_{\mathcal{IC}}^p(\mathcal{K})$, that is,

$$\text{GMod}_{\mathcal{IC}}^p(\mathcal{K}) = \{P \in \text{Mod}(\mathcal{IC}) \mid \|A_{\mathcal{K}}P\|_p = \mathcal{I}_{\mathcal{IC}}^p(\mathcal{K})\}.$$

Proposition 1. Let \mathcal{IC} be a set of integrity constraints.

1. If $\mathcal{IC} = \emptyset$, then $\mathcal{I}_{\mathcal{IC}}^p = \mathcal{I}_{\Pi}^p$ and $\text{GMod}_{\mathcal{IC}}^p(\mathcal{K}) = \text{GMod}^p(\mathcal{K})$ for all knowledge bases \mathcal{K} .
2. $\text{GMod}_{\mathcal{IC}}^p(\mathcal{K})$ is always non-empty, compact and convex.
3. If $\mathcal{K} \cup \mathcal{IC}$ is consistent, $\text{GMod}_{\mathcal{IC}}^p(\mathcal{K}) = \text{Mod}(\mathcal{K} \cup \mathcal{IC})$.

Proof sketch. 1. follows immediately from the fact that $\text{Mod}(\emptyset) = \mathcal{P}(\text{At})$ and the definitions. 2. and 3. follow exactly like the corresponding properties of $\text{GMod}^p(\mathcal{K})$ obtained in [Potyka, 2014] and [Potyka and Thimm, 2014]. \square

Now we can define the generalized entailment problem with integrity constraints.

Definition 3 (Generalized entailment problem with integrity constraints). Given a knowledge base \mathcal{K} , integrity constraints \mathcal{IC} and a query $(\psi \mid \phi)$, $\phi, \psi \in \mathcal{L}(\text{At})$, the generalized entailment problem with integrity constraints is to solve

$$\begin{aligned} \text{opt}_{P \in \text{GMod}_{\mathcal{IC}}^p(\mathcal{K})} & P(\psi \mid \phi) & (4) \\ \text{subject to} & P(\phi) > 0, \end{aligned}$$

where opt stands for min and max respectively.

We denote the generalized entailment relation by $\models_{\mathcal{IC}}^p$, i. e., if l and u are the lower and upper bounds obtained from (4), we write $\mathcal{K} \models_{\mathcal{IC}}^p (\psi \mid \phi)[l, u]$. As before, if there is no $P \in \text{GMod}_{\mathcal{IC}}^p(\mathcal{K})$ with $P(\phi) > 0$, we let $l = 1$, $u = 0$.

Before looking at this problem in more detail, we consider some examples to illustrate that generalized entailment can yield reasonable results even if \mathcal{K} is inconsistent. By reasonable we mean that the generalized entailment results can be regarded as merging contradictory opinions. The way in which the opinions are merged depends on the selected p -norm. How should we choose p ? Intuitively, $p = 1$ takes the violation of all opinions into account without regarding how strong a single opinion is violated. On the other extreme, $p = \infty$ takes only the maximal violation of a single opinion into account and ignores the overall violation of all opinions. $p = 2$ yields a good balance between both extremes.

	$\mathcal{I}_{\mathcal{IC}}^1$	$\mathcal{I}_{\mathcal{IC}}^2$	$\mathcal{I}_{\mathcal{IC}}^\infty$
\mathcal{K}_1	[0, 1]	[0.5, 0.5]	[0.5, 0.5]
\mathcal{K}_2	[0.1, 0.8]	[0.45, 0.45]	[0.45, 0.45]
\mathcal{K}_3	[0.7, 0.7]	[0.533, 0.533]	[0.45, 0.45]
\mathcal{K}_4	[0.8, 0.8]	[0.6, 0.6]	[0.5, 0.5]
\mathcal{K}_5	[0.8, 0.8]	[0.566, 0.566]	[0.45, 0.45]

Table 1: Generalized entailment results (rounded to 3 digits) for the probability of A (Example 1).

Query	$\mathcal{I}_{\mathcal{IC}}^1$	$\mathcal{I}_{\mathcal{IC}}^2$	$\mathcal{I}_{\mathcal{IC}}^\infty$
$(P \mid N)$	[0.1, 0.9]	[0.384, 0.615]	[0.376, 0.624]
$(P \mid Q)$	[0.1, 0.9]	[0.517, 0.615]	[0.520, 0.624]
$(P \mid R)$	[0.1, 0.9]	[0.384, 0.482]	[0.376, 0.481]
(N)	[1, 1]	[1, 1]	[1, 1]

Table 2: Generalized entailment results (rounded to 3 digits) for Nixon diamond with $\mathcal{IC} = \{(N)[1]\}$ (Example 2).

Example 1. Suppose we have some experts with different opinions on the probability of some event A , say, that the price of a stock rises. We consider the knowledge bases $\mathcal{K}_1 = \langle (A)[0], (A)[1] \rangle$, $\mathcal{K}_2 = \langle (A)[0.1], (A)[0.8] \rangle$, $\mathcal{K}_3 = \langle (A)[0.1], (A)[0.8], (A)[0.7] \rangle$, $\mathcal{K}_4 = \langle (A)[0.1], (A)[0.8], (A)[0.9] \rangle$, $\mathcal{K}_5 = \langle (A)[0.1], (A)[0.8], (A)[0.8] \rangle$. In \mathcal{K}_1 both experts are completely convinced of their opinion. In \mathcal{K}_2 both experts choose a more conservative formulation and in \mathcal{K}_3 , \mathcal{K}_4 and \mathcal{K}_5 we have a third expert who also thinks that A is rather likely. We do not need any integrity constraints and set $\mathcal{IC} = \emptyset$. Table 1 shows generalized entailment results for the query (A) and $p = 1, 2, \infty$.

For $p = 1$, we get most conservative results. For two experts, the whole interval between both opinions is possible. If we add a third expert, the results corresponds to the median of the experts' opinions. As p increases, larger violations are penalized more heavily and we end up with point probabilities somewhere between the experts' opinions. Finally, for $p = \infty$, only the maximal violation counts and so there is no difference between \mathcal{K}_2 , \mathcal{K}_3 and \mathcal{K}_5 since the extreme opinions are represented by probabilities 0.1 and 0.8 in each case.

Example 2. Let us consider the *Nixon diamond*. We believe that quakers (Q) are usually pacifists (P) while republicans (R) are usually not. However, we know that Nixon (N) was both a quaker and a republican. We do not doubt the existence of Nixon and therefore consider the integrity constraint $\mathcal{IC} = \langle (N)[1] \rangle$. The remaining knowledge is represented as follows: $\mathcal{K} = \langle (P \mid Q)[0.9], (P \mid R)[0.1], (QR \mid N)[1] \rangle$. Table 2 shows the generalized entailment results.

Again, $p = 1$ yields most conservative results. For $p > 1$, we maintain the knowledge that quakers are probably pacifists and that republicans are probably not.

Example 3. We consider a variant of Kyburg's *Lottery Paradox* [Kyburg, 1992] similar to [Knight, 2002]. There is a lottery and exactly one player will win. However, for a particular player p , we do not believe that p will win. We model

k	$\mathcal{I}_{\mathcal{IC}}^1$	$\mathcal{I}_{\mathcal{IC}}^2$	$\mathcal{I}_{\mathcal{IC}}^\infty$
1	[1, 1]	[1, 1]	[1, 1]
2	[0, 1]	[0.5, 0.5]	[0.5, 0.5]
4	[0, 1]	[0.25, 0.25]	[0.25, 0.25]
8	[0, 1]	[0.125, 0.125]	[0.125, 0.125]

Table 3: Probabilities that a particular player will win in the lottery paradox with k players (Example 3).

the lottery paradox with k players by the knowledge base $\mathcal{K}_k = \langle (p_1)[0], \dots, (p_k)[0] \rangle$, where p_k expresses that player k will win. The fact that one player will win is represented as an integrity constraint, i. e., $\mathcal{IC}_k = \langle (\bigvee_{i=1}^k p_i)[1] \rangle$. As the number of players goes to infinity, the degree of inconsistency of the knowledge base goes to 0. If we perform generalized entailment to compute the probability that a particular player will win, the probability is equal for each player. Table 3 shows the results for different k .

Since the knowledge base does not favor any player, we cannot conclude anything for $p = 1$. For $p > 1$, the probability that a player wins is uniformly distributed as one would expect under the given premises.

4 Analysis

Our proposal of the generalized entailment relation $\models_{\mathcal{IC}}^p$ aims at extending the classical entailment relation \models^c to inconsistent knowledge bases. The examples at the end of the previous section suggest that generalized entailment results are also intuitive in the case of inconsistency. In the spirit of other non-classical reasoning approaches like, e. g., [Kraus *et al.*, 1990], we will now propose a set of rationality postulates that each probabilistic entailment relation $\models_{\mathcal{IC}}$ should satisfy that extends probabilistic entailment to inconsistent knowledge bases while maintaining integrity constraints \mathcal{IC} .

Let \uplus denote disjoint union, i. e., $A = A_1 \uplus A_2$ means $A = A_1 \cup A_2$ and $A_1 \cap A_2 = \emptyset$. We consider the following postulates for a probabilistic entailment relation $\models_{\mathcal{IC}}$.

1. *Consistency*: If $\mathcal{K} \cup \mathcal{IC}$ is consistent, then it holds $(\mathcal{K} \cup \mathcal{IC}) \models^c (\psi | \phi)[l, u]$ iff $\mathcal{K} \models_{\mathcal{IC}} (\psi | \phi)[l, u]$.
2. *Integrity*: For all $(\psi | \phi)[d] \in \mathcal{IC}$, it holds that either $\mathcal{K} \models_{\mathcal{IC}} (\psi | \phi)[d, d]$ or $\mathcal{K} \models_{\mathcal{IC}} (\phi)[0, 0]$.
3. *Consistent Independence*: Let $\text{At} = \text{At}_1 \uplus \text{At}_2$ and let $\mathcal{K} = \mathcal{K}_1 \uplus \mathcal{K}_2$ such that \mathcal{K}_i is a knowledge base over $\mathcal{L}(\text{At}_i)$, $i = 1, 2$. If $\mathcal{K}_i \cup \mathcal{IC}$ is consistent and $(\psi_i | \phi_i)$ is a query over $\mathcal{L}(\text{At}_i)$, then $\mathcal{K}_i \cup \mathcal{IC} \models^c (\psi_i | \phi_i)[l, u]$ holds in $\mathcal{L}(\text{At}_i)$ iff $\mathcal{K} \models_{\mathcal{IC}} (\psi_i | \phi_i)[l, u]$ holds in $\mathcal{L}(\text{At})$.
4. *Independence*: Let $\text{At} = \text{At}_1 \uplus \text{At}_2$ and let $\mathcal{K} = \mathcal{K}_1 \uplus \mathcal{K}_2$ such that \mathcal{K}_i is a knowledge base over $\mathcal{L}(\text{At}_i)$, $i = 1, 2$. If $(\psi_i | \phi_i)$ is a query over $\mathcal{L}(\text{At}_i)$, then $\mathcal{K} \models_{\mathcal{IC}} (\psi_i | \phi_i)[l, u]$ holds in $\mathcal{L}(\text{At})$ if and only if $\mathcal{K}_i \models_{\mathcal{IC}} (\psi_i | \phi_i)[l, u]$ holds in $\mathcal{L}(\text{At}_i)$.
5. *Continuity*: If \mathcal{K} is 'close' to a consistent knowledge base \mathcal{K}_0 such that $(\mathcal{K}_0 \cup \mathcal{IC}) \models^c (\psi | \phi)[l_0, u_0]$, $(\mathcal{K}_0 \cup \mathcal{IC}) \not\models^c (\phi)[0, 0]$ and $\mathcal{K} \models_{\mathcal{IC}} (\psi | \phi)[l, u]$, then $[l, u]$ is 'close' to $[l_0, u_0]$.

Consistency states that the extended entailment relation should agree with probabilistic entailment if the given information is consistent. *Integrity* assures that all integrity constraints are either obeyed or not applicable at all (then $\mathcal{K} \models_{\mathcal{IC}} (\psi | \phi)[1, 0]$). *Consistent Independence* states that consistency remains true for subsets of the language if there is only consistent information about this subset. *Independence* states that knowledge about a subset of the language should not influence entailment results about the remaining language. In particular, this property can be exploited to decompose the extended entailment problem into two smaller problems. *Continuity* says that if \mathcal{K}_0 is consistent and does not classically entail $P(\phi) = 0$, then minor changes in \mathcal{K}_0 shall not result in major changes in the entailed probability $P(\psi | \phi)$ even if \mathcal{K}_0 becomes inconsistent. Since the definition of closeness is very subtle in this context, it will be discussed later on in more detail. We have the following relationships between our postulates.

Proposition 2.

1. *Consistent Independence implies Consistency.*
2. *Consistency and Independence implies Consistent Independence.*

Proof. 1 follows immediately by letting $\text{At}_2 = \mathcal{K}_2 = \emptyset$. To prove 2, note that by Independence, generalized entailment w.r.t. \mathcal{K} over $\mathcal{L}(\text{At})$ is equivalent to generalized entailment w.r.t. \mathcal{K}_1 over $\mathcal{L}(\text{At}_1)$. But since \mathcal{K}_1 is consistent, the claim follows with Consistency applied to \mathcal{K}_1 over $\mathcal{L}(\text{At}_1)$. \square

Note that there are some interesting relationships to other properties. Consistent Independence implies *Reflexivity* [Kraus *et al.*, 1990], i. e., if $(\psi | \phi)[p]$ is a satisfiable conditional and contains no atoms mentioned in \mathcal{K} , then $\mathcal{K} \cup \{(\psi | \phi)[p]\} \models_{\mathcal{IC}} (\psi | \phi)[p, p]$. Independence implies *Language Invariance* [Paris, 1994], i. e., just adding additional atoms to the language does not change the entailment results.

Generalized entailment satisfies our first four desiderata.

Theorem 1. *The generalized entailment relation $\models_{\mathcal{IC}}^p$ satisfies Consistency, Integrity, Consistent Independence and Independence.*

Proof sketch. To prove consistency, note that Proposition 1, 3, implies that $\text{GMod}_{\mathcal{IC}}^p(\mathcal{K}) = \text{Mod}(\mathcal{K} \cup \mathcal{IC})$. But then (4) is just the definition of the probabilistic entailment problem.

Integrity follows from $\text{GMod}_{\mathcal{IC}}^p(\mathcal{K}) \subseteq \text{Mod}(\mathcal{IC})$.

By consistency and Proposition 2, Consistent Independence follows from Independence.

To prove Independence, show that for each probability function P over $\Omega(\text{At})$ that satisfies \mathcal{K} , there are corresponding probability function P_i over $\Omega(\text{At}_i)$ that satisfy \mathcal{K}_i , $i = 1, 2$ and agree with P for all formulas from $\mathcal{L}(\text{At}_i)$ and vice versa. To get from P to P_i just marginalize. To get from P_1 and P_2 to P let $P(\omega) = P_1(\omega|_{\text{At}_1}) \cdot P_2(\omega|_{\text{At}_2})$. To meet space restriction, we leave out the details of the proof. \square

To prove continuity, we need a more precise notion of closeness of knowledge bases. However, using a too strong

notion of closeness, continuity cannot be satisfied by any extended entailment relation that extends probabilistic entailment in a reasonable way, because even probabilistic entailment behaves discontinuously in some cases. Consider the following non-trivial example from [Paris, 1994]¹.

Example 4. Consider a disease d , a symptom s and a possible complication c . Let \mathcal{K} contain the conditionals $(d|s)[0.75]$, $(d|\bar{s})[0.25]$, $(\bar{c}d|s)[0.15]$, $(\bar{c}|d\bar{s})[0.6]$, $(c|ds)[0.8]$ and $(cd|\bar{s})[0.1]$. \mathcal{K} is consistent and, for instance, $\mathcal{K} \models^c (\bar{s})[0, 1]$. However, if we construct \mathcal{K}' from \mathcal{K} by replacing $(cd|\bar{s})[0.1]$ with $(cd|\bar{s})[0.0999]$, we have $\mathcal{K} \models^c (\bar{s})[0, 0]$. Such discontinuities are connected to each conditional in \mathcal{K} , see [Paris, 1994], p. 90, for more details.

To exclude such discontinuities, Paris defined convergence of knowledge bases as follows: (\mathcal{K}_i) converges to \mathcal{K} iff $(\text{Mod}(\mathcal{K}_i))$ converges to $\text{Mod}(\mathcal{K})$ with respect to the Blaschke metric. Roughly speaking, $S_1, S_2 \subseteq \mathbb{R}^n$ have Blaschke distance d , $\|S_1, S_2\|_B = d$, iff for every $x_1 \in S_1$, there is a $x_2 \in S_2$ such that $\|x_1 - x_2\|_2 \leq d$ and vice versa. By replacing the models with the generalized models in this notion of convergence, we obtain the following weak form of continuity for generalized entailment.

Theorem 2 (Weak Continuity). *Let (\mathcal{K}_i) be a sequence of knowledge bases such that $(\text{GMod}_{\mathcal{IC}}^p(\mathcal{K}_i))$ converges to $\text{GMod}_{\mathcal{IC}}^p(\mathcal{K})$ with respect to the Blaschke metric. If $\mathcal{K} \not\models_{\mathcal{IC}}^p (\phi)[0, 0]$, $\mathcal{K} \models_{\mathcal{IC}}^p (\psi|\phi)[l, u]$ and $\mathcal{K}_i \models_{\mathcal{IC}}^p (\psi|\phi)[l_i, u_i]$, then l_i and u_i converge to l and u , respectively.*

Proof sketch. For ease of notation, let $G = \text{GMod}_{\mathcal{IC}}^p(\mathcal{K})$ and $G_i = \text{GMod}_{\mathcal{IC}}^p(\mathcal{K}_i)$. The claim follows from (A), where

(A) for all $\epsilon > 0$ that are sufficiently small, there is a $\delta > 0$ such that $\|G, G_i\|_B < \delta$ implies that for all $P \in G$ ($P' \in G_i$) with $P(\phi) > 0$ ($P'(\phi) > 0$) there is a $P' \in G_i$ ($P \in G$) such that $|P(\psi|\phi) - P'(\psi|\phi)| < \epsilon$

We know from Real Analysis that $\|x\|_1 \leq \sqrt{n} \|x\|_2$ for all $x \in \mathbb{R}^n$. Note also that $|P(F) - P'(F)| \leq \|P - P'\|_1$ for all $F \in \mathcal{L}(\text{At})$. Therefore, $\|G, G_i\|_B < \epsilon/\sqrt{n}$ implies that for all $P \in G$ ($P' \in G_i$), there is a $P' \in G_i$ ($P \in G$) with $|P(F) - P'(F)| < \epsilon$. Since $\mathcal{K} \not\models_{\mathcal{IC}}^p (\phi)[0, 0]$, there is a $P \in G$ with $P(\phi) > 0$. Hence, if $\delta < \frac{P(\phi)}{2\sqrt{n}}$, there is a $P' \in G_i$ with $P'(\phi) > P(\phi)/2 > 0$. Hence, if δ is sufficiently small, both $[l, u]$ and $[l_i, u_i]$ are non-trivial. Finally, check that for $0 < \epsilon < 1$ and $\delta < \frac{\epsilon P(\phi)}{4\sqrt{n}}$ ($\delta < \frac{\epsilon P_i(\phi)}{4\sqrt{n}}$), (A) holds. \square

Note that if $(\mathcal{K} \cup \mathcal{IC})$ is consistent, Consistency implies that $(\mathcal{K} \cup \mathcal{IC}) \not\models_{\mathcal{IC}}^p (\phi)[0, 0]$ and $(\mathcal{K} \cup \mathcal{IC}) \models^c (\psi|\phi)[l, u]$, so that l_i and u_i converge to the probabilistic entailment result as demanded in Convergence.

5 Computational Aspects

We cannot expect to find highly efficient algorithms for the generalized entailment problem, since even the probabilistic satisfiability problem is NP-hard [Georgakopoulos *et al.*,

1988]. However, it is interesting to ask how much more difficult is the generalized entailment problem as compared to the probabilistic entailment problem.

Our first goal is to show that the generalized entailment problem can be solved by linear programming techniques. To do so, we introduce a vector $a_F = (1_{\{F\}}(\omega_j))_{1 \leq j \leq n}$ for each formula F . Note that $a_F P = P(F)$, see also [Nilsson, 1986; Jaumard *et al.*, 1991]. We will also need the following lemma, which is a straightforward generalization of [Potyka and Thimm, 2014], Lemma 1.

Lemma 1. *Let \mathcal{K} be a knowledge base, let \mathcal{IC} be a set of integrity constraints and let $1 < p < \infty$. Let $P \in \text{GMod}_{\mathcal{IC}}^p(\mathcal{K})$ be a generalized model and let $x = A_{\mathcal{K}} P$. Then it holds $A_{\mathcal{K}} P' = x$ for all $P' \in \text{GMod}_{\mathcal{IC}}^p(\mathcal{K})$ and we call $x = x_{\mathcal{K}}^p$ the violation vector of \mathcal{K} .*

Theorem 3. *The generalized entailment problem with integrity constraints has a well-defined solution and (4) is equivalent to the following linear programs, where $\epsilon_p = \mathcal{I}_{\mathcal{IC}}^p(\mathcal{K})$, \mathbb{R}_+^n denotes the non-negative real vectors and opt stands for \min and \max , respectively.*

- For $p = 1$, (4) is equivalent to

$$\begin{aligned} & \text{opt}_{(x,y,t) \in \mathbb{R}_+^{n+m+1}} && a_{\psi\phi} x && (5) \\ & \text{subject to} && -y \leq A_{\mathcal{K}} x \leq y, \\ & && \sum_{i=1}^m y_i = t \cdot \epsilon_1, A_{\mathcal{IC}} x = 0, a_{\top} x = t, a_{\phi} x = 1. \end{aligned}$$

- For $1 < p < \infty$, (4) is equivalent to

$$\begin{aligned} & \text{opt}_{(x,t) \in \mathbb{R}_+^{n+1}} && a_{\psi\phi} x && (6) \\ & \text{subject to} && A_{\mathcal{K}} x = t \cdot \epsilon_p, A_{\mathcal{IC}} x = 0, a_{\top} x = t, a_{\phi} x = 1. \end{aligned}$$

- For $p = \infty$, (4) is equivalent to

$$\begin{aligned} & \text{opt}_{(x,t) \in \mathbb{R}_+^{n+1}} && a_{\psi\phi} x && (7) \\ & \text{subject to} && -t \cdot \epsilon_{\infty} \leq A_{\mathcal{K}} x \leq t \cdot \epsilon_{\infty}, \\ & && A_{\mathcal{IC}} x = 0, a_{\top} x = t, a_{\phi} x = 1. \end{aligned}$$

In particular, the linear programs are feasible if and only if there is a $P \in \text{GMod}_{\mathcal{IC}}^p(\mathcal{K})$ with $P(\phi) > 0$.

Proof sketch. To begin with, recall that $\text{GMod}_{\mathcal{IC}}^p(\mathcal{K}) \neq \emptyset$ is guaranteed. (4) can be rewritten as

$$\text{opt}_{x \in \mathbb{R}_+^n} \frac{a_{\psi\phi} x}{a_{\phi} x} \quad (8)$$

subject to

$$\|A_{\mathcal{K}} x\|_p = \epsilon_p, A_{\mathcal{IC}} x = 0, a_{\top} x = 1, a_{\phi} x > 0,$$

To get rid of the non-linear constraint $\|A_{\mathcal{K}} x\|_p = \epsilon_p$, we can apply Lemma 1 for $1 < p < \infty$, to replace $\|A_{\mathcal{K}} x\|_p = \epsilon_p$ with $A_{\mathcal{K}} P' = x_{\mathcal{K}}^p$. For $p = 1$ and $p = \infty$, we can exploit piecewise linearity to replace $\|A_{\mathcal{K}} x\|_1 = \epsilon_1$ with the constraints $-y \leq A_{\mathcal{K}} x \leq y$ and $\sum_{i=1}^m y_i = \epsilon_1$, where $y \in \mathbb{R}^m$; and to replace $\|A_{\mathcal{K}} x\|_{\infty} = \epsilon_{\infty}$ with $-\epsilon_{\infty} \leq A_{\mathcal{K}} x \leq \epsilon_{\infty}$.

¹The example was originally proposed in P. Courtney, Doctoral thesis, Manchester University, Manchester, U.K., 1992.

Problem	\bar{n}	\bar{m}	Cost
PSAT	$ \Omega + \mathcal{K} $	$ \mathcal{K} $	$\bar{n} \bar{m}^2$
$\mathcal{I}_{\mathcal{IC}}^1$	$ \Omega + 3 \mathcal{K} $	$2 \mathcal{K} + \mathcal{IC} $	$\bar{n} \bar{m}^2$
$\mathcal{I}_{\mathcal{IC}}^2$	$ \Omega $	$ \mathcal{IC} $	$\bar{n}^2 \bar{m}$
$\mathcal{I}_{\mathcal{IC}}^\infty$	$ \Omega + 2 \mathcal{K} $	$2 \mathcal{K} + \mathcal{IC} $	$\bar{n} \bar{m}^2$
$\mathcal{I}_{\mathcal{IC}}^p$	$ \Omega $	$ \mathcal{IC} $	\bar{n}^3
PENT	$ \Omega $	$ \mathcal{K} $	$\bar{n} \bar{m}^2$
GENT $p = 1$	$ \Omega + 3 \mathcal{K} $	$2 \mathcal{K} + \mathcal{IC} $	$\bar{n} \bar{m}^2$
GENT $1 < p < \infty$	$ \Omega $	$ \mathcal{K} + \mathcal{IC} $	$\bar{n} \bar{m}^2$
GENT $p = \infty$	$ \Omega + 2 \mathcal{K} $	$2 \mathcal{K} + \mathcal{IC} $	$\bar{n} \bar{m}^2$

Table 4: Number of optimization variables \bar{n} , number of constraints \bar{m} (ignoring constants and non-negativity constraints) and rough performance estimates for testing satisfiability (PSAT), computing minimal violations measures, probabilistic entailment (PENT) and generalized entailment (GENT) with standard algorithms.

To get rid of the non-linear objective, we can apply a result from [Charnes and Cooper, 1962], which is also used to solve the probabilistic entailment problem [Jaumard *et al.*, 1991]. Basically, the feasible solutions are scaled such that $a_\phi x = 1$ holds. Then $\frac{a_\psi \phi x}{a_\phi x}$ equals $a_\psi \phi x$. This transformation does not change the optimal objective as the scaling factor cancels out in the fraction, see [Charnes and Cooper, 1962] for details.

Equivalence of the linear programs with (2) follows with the arguments sketched above and guarantees that all linear programs are feasible if and only if $P(\phi) > 0$ for some $P \in \text{GMod}_{\mathcal{IC}}^p(\mathcal{K})$. If there is some $P \in \text{GMod}_{\mathcal{IC}}^p(\mathcal{K})$ with $P(\phi) > 0$, existence of the solutions follows from the theory of linear programming. \square

Now let us look at the cost of solving the generalized entailment problem. Reasoning is usually a two-stage process. First, we test satisfiability, then we perform a reasoning algorithm. In our approach, the satisfiability phase is replaced with an inconsistency measuring phase. To compute minimal violation measures, we have to solve a linear program for $p \in \{1, \infty\}$, a quadratic program for $p = 2$ and a convex program for other p , see [Potyka, 2014] for details. Probabilistic satisfiability, probabilistic entailment and generalized entailment can be computed by linear programs. Expected costs when using standard algorithms are summarized in Table 4. For linear programs, we consider estimates proposed in [Matousek and Gärtner, 2007] for the Simplex algorithm. For quadratic and convex programs, we use estimates proposed in [Boyd and Vandenberghe, 2004] for interior-point methods. The cost is estimated with respect to the number of optimization variables \bar{n} and the number of constraints \bar{m} . Note that we have to introduce additional slack variables for linear programs whenever inequalities are present.

In practice, $|\Omega|$ is the dominating factor because it depends exponentially on the number of propositional variables $|\text{At}|$ in our language. In contrast, $|\mathcal{K}|$ usually grows at most polynomially in $|\text{At}|$. In fact, whenever $|\mathcal{K}| > |\Omega|$ the knowledge base is overdetermined in the sense that it either contains redundant information ($A_{\mathcal{K}}$ does not have full rank) or

it leaves no degrees of freedom ($A_{\mathcal{K}}x$ has at most one solution). Taking this into account, we see that computing minimal violation measures for $p = 1$ and $p = \infty$ is asymptotically not harder than testing satisfiability. Similarly, performing generalized entailment, when the inconsistency values are known, is asymptotically not harder than performing probabilistic entailment. In fact, for $1 < p < \infty$, we get basically the same cost because the violation constraints can be represented by linear equalities as explained in Lemma 1. To deal with larger instances of the generalized entailment problem, we can exploit *Independence* and apply column generation techniques to reduce the exponential influence of $|\text{At}|$ on $|\Omega|$ [Hansen and Perron, 2008; Finger and De Bona, 2011; Cozman and Ianni, 2013].

6 Related Work

An overview of inconsistency measures for classical logics can be found in [Grant and Hunter, 2013b], an overview of measures for probabilistic logics in [Thimm, 2013]. The idea of generalized reasoning transfers primarily to approaches that measure inconsistency by a notion of distance from interpretations to actual models. An interesting family of such measures for classical logics has been proposed in [Grant and Hunter, 2013a]. The idea is to extend the models of single formulas in the knowledge base until the intersection for all formulas is non-empty. The resulting set can be understood as a classical notion of a set of generalized models and it is interesting to ask if reasonable generalized inference relations for classical logics can be derived. Note also that minimal violation measures have recently been generalized to languages allowing probability intervals $[l, u]$, $0 \leq l \leq u \leq 1$ rather than point probabilities $[d]$ and some properties have been strengthened in this framework [De Bona and Finger, 2014].

To deal with inconsistencies in classical logics, several approaches have been proposed. For instance, one can introduce new connectives, consider consistent subsets of the knowledge base or apply belief merging approaches [Koniczny *et al.*, 2005; Béziau *et al.*, 2007; Koniczny and Perez, 2011]. For probabilistic logics, several revision, fusion and merging approaches have been considered, see, for instance, [Kern-Isberner and Rödder, 2004; Weydert, 2011; Wilmers, 2015].

The idea of generalizing the notion of a probabilistic model has also been employed in [Daniel, 2009]. There, reasoning in inconsistent probabilistic knowledge is realized by a fuzzy notion of a model and this is used to generalize reasoning based on the principle of maximum entropy. However, the general probabilistic entailment problem and computational issues are not discussed in [Daniel, 2009].

7 Summary

We defined the generalized entailment problem with integrity constraints and showed that it satisfies several desirable properties. These properties seem to be reasonable desiderata for each approach that extends probabilistic entailment to inconsistent knowledge bases. Generalized entailment satisfies only a weak form of continuity, but this seems to be true for all reasonable extensions because of discontinuities that are

inherent to the probabilistic entailment problem. Computationally, generalized entailment for $p = 1, \infty$ is barely harder than performing a probabilistic satisfiability test and probabilistic entailment.

The approach proposed in this paper has been implemented in Java and is available as open source².

References

- [Bench-Capon and Dunne, 2007] T. J. M. Bench-Capon and P. E. Dunne. Argumentation in Artificial Intelligence. *Artificial Intelligence*, 171:619–641, 2007.
- [Benferhat *et al.*, 1999] S. Benferhat, D. Dubois, and H. Prade. Possibilistic and standard probabilistic semantics of conditional knowledge bases. *Journal of Logic and Computation*, 9(6):873–895, 1999.
- [Béziau *et al.*, 2007] J.-Y. Béziau, W. Carnielli, and D. Gabbay, editors. *Handbook of Paraconsistency*. College Publications, London, 2007.
- [Boyd and Vandenberghe, 2004] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [Charnes and Cooper, 1962] A. Charnes and W. W. Cooper. Programming with linear fractional functionals. *Naval Research logistics quarterly*, 9(3-4):181–186, 1962.
- [Cozman and Ianni, 2013] F. G. Cozman and L. F. Ianni. Probabilistic satisfiability and coherence checking through integer programming. In *Proceedings ECSQARU2013*, pages 145–156. Springer, 2013.
- [Daniel, 2009] L. Daniel. *Paraconsistent Probabilistic Reasoning*. PhD thesis, L'École Nationale Supérieure des Mines de Paris, 2009.
- [De Bona and Finger, 2014] G. De Bona and M. Finger. Notes on measuring inconsistency in probabilistic logic. Technical report, RT-MAC-2014-02, IME/USP, 2014.
- [Finger and De Bona, 2011] M. Finger and G. De Bona. Probabilistic satisfiability: Logic-based algorithms and phase transition. In *Proceedings IJCAI 2011*, pages 528–533, 2011.
- [Georgakopoulos *et al.*, 1988] G. F. Georgakopoulos, D. J. Kavvadias, and C. H. Papadimitriou. Probabilistic satisfiability. *Journal of Complexity*, 4(1):1–11, 1988.
- [Grant and Hunter, 2013a] J. Grant and A. Hunter. Distance-based measures of inconsistency. In *Proceedings ECSQARU2013*, pages 230–241, 2013.
- [Grant and Hunter, 2013b] J. Grant and A. Hunter. Measuring the good and the bad in inconsistent information. In *Proceedings IJCAI 2013*, 2013.
- [Hansen and Perron, 2008] P. Hansen and S. Perron. Merging the local and global approaches to probabilistic satisfiability. *International Journal of Approximate Reasoning*, 47(2):125–140, 2008.
- [Hansson, 2001] S. O. Hansson. *A Textbook of Belief Dynamics*. Kluwer Academic Publishers, 2001.
- [Hunter and Konieczny, 2010] A. Hunter and S. Konieczny. On the measure of conflicts: Shapley inconsistency values. *Artificial Intelligence*, 174(14):1007–1026, 2010.
- [Jaumard *et al.*, 1991] B. Jaumard, P. Hansen, and Marcus Poggi. Column generation methods for probabilistic logic. *ORSA - Journal on Computing*, 3(2):135–148, 1991.
- [Kern-Isberner and Rödder, 2004] G. Kern-Isberner and W. Rödder. Belief revision and information fusion on optimum entropy. *International Journal of Intelligent Systems*, 19(9):837–857, 2004.
- [Kern-Isberner, 2001] G. Kern-Isberner. *Conditionals in Nonmonotonic Reasoning and Belief Revision*. Springer, 2001.
- [Knight, 2002] K. Knight. Measuring inconsistency. *Journal of Philosophical Logic*, 31(1):77–98, 2002.
- [Konieczny and Perez, 2011] S. Konieczny and R. Pino Perez. Logic based merging. *Journal of Philosophical Logic*, 40:239–270, 2011.
- [Konieczny *et al.*, 2005] Sébastien Konieczny, Jérôme Lang, and Pierre Marquis. Reasoning under inconsistency: the forgotten connective. In *Proceedings IJCAI 2005*, pages 484–489, 2005.
- [Kraus *et al.*, 1990] S. Kraus, D. J. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44(1-2):167–207, 1990.
- [Kyburg, 1992] H. Kyburg. *Probability and the logic of rational belief*. UMI, 1992.
- [Lukasiewicz, 1999] Thomas Lukasiewicz. Probabilistic deduction with conditional constraints over basic events. *JAIR*, 10:380–391, 1999.
- [Matousek and Gärtner, 2007] J. Matousek and B. Gärtner. *Understanding and Using Linear Programming*. Universitext (1979). Springer, 2007.
- [Nilsson, 1986] N. J. Nilsson. Probabilistic logic. *Artificial Intelligence*, 28:71–88, 1986.
- [Paris, 1994] J. B. Paris. *The Uncertain Reasoner's Companion – A Mathematical Perspective*. Cambridge University Press, 1994.
- [Picado-Muiño, 2011] D. Picado-Muiño. Measuring and Repairing Inconsistency in Probabilistic Knowledge Bases. *International Journal of Approximate Reasoning*, 2011.
- [Potyka and Thimm, 2014] N. Potyka and M. Thimm. Consolidation of probabilistic knowledge bases by inconsistency minimization. In *Proceedings ECAI 2014*, pages 729–734. IOS Press, 2014.
- [Potyka, 2014] N. Potyka. Linear programs for measuring inconsistency in probabilistic logics. In *Proceedings KR 2014*. AAAI Press, 2014.
- [Reiter, 1980] R. Reiter. A logic for default reasoning. *Artificial Intelligence*, 13:81–132, 1980.

²tweetyproject.org

- [Thimm, 2013] M. Thimm. Inconsistency measures for probabilistic logics. *Artificial Intelligence*, 197:1–24, April 2013.
- [Weydert, 2011] E. Weydert. Conditional ranking revision - iterated revision with sets of conditionals. *Journal of Philosophical Logic*, 41(9):237–271, 2011.
- [Wilmers, 2015] George Wilmers. A foundational approach to generalising the maximum entropy inference process to the multi-agent context. *Entropy*, 17(2):594–645, 2015.