

On Partial Information and Contradictions in Probabilistic Abstract Argumentation

Anthony Hunter

University College London, UK
anthony.hunter@ucl.ac.uk

Matthias Thimm

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

Abstract

We provide new insights into the area of combining abstract argumentation frameworks with probabilistic reasoning. In particular, we consider the scenario when assessments on the probabilities of a subset of the arguments is given and the probabilities of the remaining arguments have to be derived, taking both the topology of the argumentation framework and principles of probabilistic reasoning into account. We generalize this scenario by also considering inconsistent assessments, i. e., assessments that contradict the topology of the argumentation framework. Building on approaches to inconsistency measurement, we present a general framework to measure the amount of conflict of these assessments and provide a method for inconsistent-tolerant reasoning.

1 Introduction

Combining approaches to qualitative and quantitative uncertain reasoning seems to be a natural way to benefit from the advantages of both areas. In this paper, we address the challenge of combining abstract argumentation frameworks (Dung 1995) with probabilistic reasoning capabilities, which has recently gained some attention in the community of formal argumentation (Dung and Thang 2010; Li, Oren, and Norman 2011; Hunter 2012; Rienstra 2012; Thimm 2012; Hunter 2013; Fazzinga, Flesca, and Parisi 2013; Li, Oren, and Norman 2013; Verheij 2014; Dondio 2014; Baroni, Giacomin, and Vicig 2014; Polberg and Doder 2014; Doder and Woltran 2014; Hunter 2014b; Hunter and Thimm 2014c; 2014b; 2014d; 2014a; Hunter 2014a; Timmer et al. 2015; Hunter 2015; Hadoux et al. 2015; Gabbay and Rodrigues 2015). An abstract argumentation framework is a directed graph with the nodes being the arguments and edges indicating attack between arguments. Work in this field w. r. t. probabilistic reasoning can be divided (Hunter 2013) into the *constellations* approach (see e. g. (Li, Oren, and Norman 2011)) and the *epistemic* approach (see e. g. (Thimm 2012)).

In the constellations approach, uncertainty in the topology of the graph (probabilities on arguments and attacks) is used to make probabilistic assessments on the acceptance of arguments. In the epistemic approach, the topology of the graph is fixed but probabilistic assessments on the acceptance of

arguments are evaluated w. r. t. the relations of the arguments in the graph. The core idea of the epistemic approach is that the more likely it is to believe in an argument, the less likely it is to believe in an argument attacking it.

The epistemic approach is useful for modeling the belief that an opponent might have in the arguments that could be presented, which is useful for example when deciding on the best arguments to present in order to persuade that opponent, see e. g. (Hunter 2015). The approach is also useful for modeling agents who are unable to directly add or change the argument graph, for instance when considering the beliefs of the audience of a debate.

Here we follow the epistemic approach to probabilistic argumentation and continue previous works (Thimm 2012; Hunter 2014a; Hunter and Thimm 2014c; 2014b; 2014d). In (Hunter and Thimm 2014d) we presented rationality conditions for epistemic probabilities on arguments based on the topology of the argumentation graph that describe when a probability assessment on arguments complies with the attack relation in an argumentation framework. In the present paper, we consider the case when these probability assessments are either incomplete or contradictory (or both) and the challenge is to complete and consolidate them. This has not been considered in the works cited above so far, with the exception of our previous work (Hunter and Thimm 2014d) where we initiated this discussion.

The central challenge in our investigation is, given probabilistic assessments on arguments that are not meaningful w. r. t. the rationality conditions of (Hunter and Thimm 2014d), how can these probabilities be modified to comply with these conditions? For this purpose and motivated by similar approaches to inconsistency measurement for classical and probabilistic logics (Hunter and Konieczny 2010; Thimm 2013; De Bona and Finger 2015), we present inconsistency measures for evaluating the appropriateness of (partial) probability assessments and a general approach to use those measures to consolidate these assessments.

The contributions of this paper are as follows.

1. We present the concept of partial probability assessments and an approach to complete them by using the principle of maximum entropy (Section 3).
2. We introduce inconsistency measures for evaluating the significance of a partial probability assessment violating

the rationality conditions (Section 4.1).

3. We use the inconsistency measures to define consolidation operators for partial probability assessments (Sections 4.2 and 4.3).

Furthermore, we provide some necessary preliminaries in Section 2, discuss related works in Section 5, and conclude with a discussion in Section 6. Proofs of technical results can be found in an online appendix¹.

2 Preliminaries

In the following, we introduce the background on abstract argumentation (Dung 1995) and the epistemic approach to probabilistic argumentation (Hunter and Thimm 2014d).

2.1 Abstract Argumentation

An *abstract argumentation framework* AF is a tuple $AF = (\text{Arg}, \rightarrow)$ where Arg is a set of arguments and \rightarrow is a relation $\rightarrow \subseteq \text{Arg} \times \text{Arg}$. For two arguments $\mathcal{A}, \mathcal{B} \in \text{Arg}$ the relation $\mathcal{A} \rightarrow \mathcal{B}$ means that argument \mathcal{A} attacks argument \mathcal{B} . For $\mathcal{A} \in \text{Arg}$ define $\mathcal{A}^- = \{\mathcal{B} \mid \mathcal{B} \rightarrow \mathcal{A}\}$.

Semantics are given to abstract argumentation frameworks by means of extensions (Dung 1995) or labelings (Wu and Caminada 2010). In this work, we use the latter. A labeling L is a function $L : \text{Arg} \rightarrow \{\text{in}, \text{out}, \text{undec}\}$ that assigns to each argument $\mathcal{A} \in \text{Arg}$ either the value in, meaning that the argument is accepted, out, meaning that the argument is not accepted, or undec, meaning that the status of the argument is undecided. Let $\text{in}(L) = \{\mathcal{A} \mid L(\mathcal{A}) = \text{in}\}$ and $\text{out}(L)$ resp. $\text{undec}(L)$ be defined analogously. A labeling L is called *conflict-free* if for no $\mathcal{A}, \mathcal{B} \in \text{in}(L)$, $\mathcal{A} \rightarrow \mathcal{B}$.

Arguably, the most important property of a semantics is its admissibility. A labeling L is called *admissible* if and only if for all arguments $\mathcal{A} \in \text{Arg}$

1. if $L(\mathcal{A}) = \text{out}$ then there is $\mathcal{B} \in \text{Arg}$ with $L(\mathcal{B}) = \text{in}$ and $\mathcal{B} \rightarrow \mathcal{A}$, and
2. if $L(\mathcal{A}) = \text{in}$ then $L(\mathcal{B}) = \text{out}$ for all $\mathcal{B} \in \text{Arg}$ with $\mathcal{B} \rightarrow \mathcal{A}$,

and it is called *complete* if, additionally, it satisfies

3. if $L(\mathcal{A}) = \text{undec}$ then there is no $\mathcal{B} \in \text{Arg}$ with $\mathcal{B} \rightarrow \mathcal{A}$ and $L(\mathcal{B}) = \text{in}$ and there is a $\mathcal{B}' \in \text{Arg}$ with $\mathcal{B}' \rightarrow \mathcal{A}$ and $L(\mathcal{B}') \neq \text{out}$.

The intuition behind admissibility is that an argument can only be accepted if there are no attackers that are accepted and if an argument is not accepted then there has to be some reasonable grounds. The idea behind the completeness property is that the status of an argument is only undec if it cannot be classified as in or out. Different types of classical semantics can be phrased by imposing further constraints. In particular, a complete labelling L is *grounded* if and only if $\text{in}(L)$ is minimal, it is *preferred* if and only if $\text{in}(L)$ is maximal, and it is *stable* if and only if $\text{undec}(L) = \emptyset$ (all statements on minimality/maximality are meant to be with respect to set inclusion).

2.2 Probabilistic Abstract Argumentation

Let $2^{\mathcal{X}}$ denote the power set of a set \mathcal{X} . A *probability function* P on some finite set \mathcal{X} is a function $P : 2^{\mathcal{X}} \rightarrow [0, 1]$ with $\sum_{X \subseteq \mathcal{X}} P(X) = 1$. Let $P : 2^{\text{Arg}} \rightarrow [0, 1]$ be a probability function on Arg . We abbreviate

$$P(\mathcal{A}) = \sum_{\mathcal{A} \in E \subseteq \text{Arg}} P(E)$$

This means that the probability of an argument is the sum of the probabilities of all sets of arguments that contain that argument. A probability function P on Arg can be regarded as a generalization of a labelling (Thimm 2012) where probability 1 suggests strong acceptance of an argument, 0 means strong rejection of an argument, and values strictly between 0 and 1 model tendencies between acceptance and rejection while probability 0.5 models undecidedness. In order to assess whether a probability function is meaningful in the context of an argumentation framework, in (Hunter and Thimm 2014d) several rationality conditions were presented. Some of them are the following:

RAT P is *rational* w.r.t. AF if for every $\mathcal{A}, \mathcal{B} \in \text{Arg}$, $\mathcal{A} \rightarrow \mathcal{B}$ then $P(\mathcal{A}) > 0.5$ implies $P(\mathcal{B}) \leq 0.5$.

COH P is *coherent* w.r.t. AF if for every $\mathcal{A}, \mathcal{B} \in \text{Arg}$, $\mathcal{A} \rightarrow \mathcal{B}$ then $P(\mathcal{A}) \leq 1 - P(\mathcal{B})$.

SFOU P is *semi-founded* w.r.t. AF if $P(\mathcal{A}) \geq 0.5$ for every $\mathcal{A} \in \text{Arg}$ with $\mathcal{A}^- = \emptyset$.

FOU P is *founded* w.r.t. AF if $P(\mathcal{A}) = 1$ for every $\mathcal{A} \in \text{Arg}$ with $\mathcal{A}^- = \emptyset$.

SOPT P is *semi-optimistic* w.r.t. AF if $P(\mathcal{A}) \geq 1 - \sum_{\mathcal{B} \in \mathcal{A}^-} P(\mathcal{B})$ for every $\mathcal{A} \in \text{Arg}$ with $\mathcal{A}^- \neq \emptyset$.

OPT P is *optimistic* w.r.t. AF if $P(\mathcal{A}) \geq 1 - \sum_{\mathcal{B} \in \mathcal{A}^-} P(\mathcal{B})$ for every $\mathcal{A} \in \text{Arg}$.

JUS P is *justifiable* w.r.t. AF if P is coherent and optimistic.

RAT ensures that if argument \mathcal{A} attacks argument \mathcal{B} , and \mathcal{A} is believed (i. e. $P(\mathcal{A}) > 0.5$), then \mathcal{B} is not believed (i. e. $P(\mathcal{B}) \leq 0.5$); COH ensures that if argument \mathcal{A} attacks argument \mathcal{B} , then the degree to which \mathcal{A} is believed caps the degree to which \mathcal{B} can be believed; SFOU ensures that if an argument is not attacked, then the argument is not disbelieved (i. e. $P(\mathcal{A}) \geq 0.5$); FOU ensures that if an argument is not attacked, then the argument is believed without doubt (i. e. $P(\mathcal{A}) = 1$); SOPT ensures that the belief in \mathcal{A} is bounded from below if the belief in its attackers is not high; OPT ensures that if an argument is not attacked, then the argument is believed without doubt (i. e. $P(\mathcal{A}) = 1$) and that the belief in \mathcal{A} is bounded from below if the belief in its attackers is not high; and JUS combines COH and OPT to provide bounds on the belief in an argument based on the belief in its attackers and attackees. We refer the reader to (Thimm 2012; Hunter 2013; Hunter and Thimm 2014c) for some more detailed discussion of these rationality conditions.

Let \mathcal{P} be the set of all probability functions, $\mathcal{P}(\text{AF})$ be the set of all probability functions on Arg , and $\mathcal{P}_t(\text{AF})$ be the set of all t -probability functions with $t \in \{\text{RAT}, \text{COH}, \text{SFOU}, \text{FOU}, \text{OPT}, \text{SOPT}, \text{JUS}\}$.

¹ mthimm.de/misc/probarg_kr2016_proofs.pdf

	\mathcal{A}_1	\mathcal{A}_2	\mathcal{A}_3	\mathcal{A}_4	\mathcal{A}_5	\mathcal{A}_6
P_1	0.7	0.4	0.8	0.3	0.5	0.6
P_2	0.6	0.3	0.6	0.3	0.4	0.6
P_3	0.2	0.7	0.6	0.3	0.6	1
P_4	0.7	0.3	0.5	0.5	0.2	0.5
P_5	0.7	0.3	0.7	0.3	0	1
P_6	0.7	0.8	0.9	0.8	0.7	1

Table 1: Some probability functions for Example 1

Example 1.] Consider the abstract argumentation framework $\text{AF} = (\text{Arg}, \rightarrow)$ depicted in Fig. 1 and the probability functions there are only partially defined by giving the probabilities of arguments). The following observations can be made:

- P_1 is rational and semi-founded, but neither coherent, founded, optimistic, semi-optimistic, nor justifiable.
- P_2 is rational, semi-founded, and coherent, but neither founded, optimistic, semi-optimistic, nor justifiable.
- P_3 is semi-founded, founded, but neither rational, coherent, optimistic, semi-optimistic, nor justifiable.
- P_4 is rational, coherent, and semi-optimistic, but neither semi-founded, founded, optimistic, nor justifiable.
- P_5 is rational, coherent, semi-optimistic, semi-founded, founded, optimistic, and justifiable.
- P_6 is semi-founded, founded, optimistic, and semi-optimistic, but neither rational, coherent, nor justifiable.




Figure 1: The argumentation framework from Example 1

A labeling L and a probability function P are *congruent*, denoted by $L \sim P$, if for all $\mathcal{A} \in \text{Arg}$ we have $L(\mathcal{A}) = \text{in} \Leftrightarrow P(\mathcal{A}) = 1$, $L(\mathcal{A}) = \text{out} \Leftrightarrow P(\mathcal{A}) = 0$, and $L(\mathcal{A}) = \text{undec} \Leftrightarrow P(\mathcal{A}) = 0.5$. The *entropy* $H(P)$ of a probability function P (Jaynes 1957; Paris 1994) is defined via

$$H(P) = - \sum_{E \subseteq \text{Arg}} P(E) \log P(E)$$

With these notations some relationships between probabilistic and classical abstract argumentation can be identified as follows.

Theorem 1. Let $P \in \mathcal{P}(\text{AF})$ and L a labelling with $L \sim P$.

1. If L is admissible then $P \in \mathcal{P}_{\text{SFOU}}(\text{AF})$.




Figure 2: The argumentation framework from Example 2

2. If L is complete then $P \in \mathcal{P}_{\text{COH}}(\text{AF}) \cap \mathcal{P}_{\text{FOU}}(\text{AF})$.
3. L is grounded iff $\{P\} = \arg \max_{Q \in \mathcal{P}_{\text{JUS}}(\text{Arg})} H(Q)$.
4. If L is stable then $P \in \arg \min_{Q \in \mathcal{P}_{\text{JUS}}(\text{Arg})} H(Q)$.

For the proofs of statements 2–4 see (Thimm 2012), the proof of statement 1 is straightforward. The uniquely defined probability function with maximum entropy (see item 3 in Theorem 1) has also a specific meaning in probabilistic reasoning. Given a set of probability functions which are compliant with a given set of constraints (the explicit knowledge in a knowledge base), the probability function with maximum entropy P_{ME} completes this knowledge in the most unbiased manner possible. That is, P_{ME} represents the explicit knowledge but adds as little knowledge as necessary in order to obtain a complete probability function. It can be characterized as the only approach satisfying some very simple rationality conditions (such as syntax insensitivity and decomposability on disjunct sub-bases), see e. g. (Paris 1994; Kern-Isberner 2001) for some details.

3 Partial Probability Assessments

The framework outlined so far allows us to assess whether probability functions are somewhat adequately reflecting the topology of an argumentation framework and, thus, can be used for uncertain reasoning based on argumentation. In the following, we will investigate the case when we already have probabilistic information on some arguments and need to infer meaningful probabilities for the remaining arguments.

Example 2. Consider a court case where the defendant John is either innocent (\mathcal{I}) or guilty (\mathcal{G}) to have committed the murder of Frank. Footage from a surveillance camera at the crime scene (\mathcal{S}_1) gives evidence that a person looking like John was present at the time of the crime, giving a reason that John is not innocent. However, footage from another surveillance camera far away from the crime scene (\mathcal{S}_2) gives evidence that a person looking like John was not present at the time of the crime, giving a reason that John is not guilty. This scenario can be modeled with the argumentation framework depicted in Figure 2. Note that there is no attack from \mathcal{G} to \mathcal{I} as a person is assumed to be innocent by default and \mathcal{I} has to be defeated explicitly to prove guilt.

Now the footage from the camera \mathcal{S}_1 is examined by a lab which assesses that the probability of the person in the pictures is indeed John is 0.7. So given $P(\mathcal{S}_1) = 0.7$ what are now adequate probabilities for the remaining arguments?

A partial function $\beta : \text{Arg} \rightarrow [0, 1]$ on Arg is called a *partial probability assignment*. Let $\text{dom } \beta \subseteq \text{Arg}$ be the domain of β , i. e., the arguments for which a probabilistic

assessment is available. We are now interested in deriving probabilities for the remaining arguments $\text{Arg} \setminus \text{dom } \beta$, taking the information we already have in β and the argumentation framework AF into account.

A probability function $P \in \mathcal{P}(\text{AF})$ is β -compliant if for every $\mathcal{A} \in \text{dom } \beta$ we have $\beta(\mathcal{A}) = P(\mathcal{A})$. Let $\mathcal{P}^\beta(\text{AF}) \subseteq \mathcal{P}(\text{AF})$ be the set of all β -compliant probability functions. Observe that $\mathcal{P}^\beta(\text{AF})$ is always non-empty.

Proposition 1. For all partial $\beta : \text{Arg} \rightarrow [0, 1]$, $\mathcal{P}^\beta(\text{AF}) \neq \emptyset$.

Of course, not all probability functions in $\mathcal{P}^\beta(\text{AF})$ are adequate for reasoning as they may not take the actual argumentation framework into account. Let $T \subseteq \{\text{RAT}, \text{COH}, \text{SFOU}, \text{FOU}, \text{OPT}, \text{SOPT}, \text{JUS}\}$ be a set of rationality conditions we wish to take into account for probabilistic reasoning on argumentation frameworks and define

$$\mathcal{P}_T(\text{AF}) = \bigcap_{t \in T} \mathcal{P}_t(\text{AF})$$

The set $\mathcal{P}_T(\text{AF})$ contains all probability functions which comply with all considered rationality conditions. Given a partial probability assessment β and some rationality conditions T , for the remainder of this section we assume $\mathcal{P}_T(\text{AF}) \cap \mathcal{P}^\beta(\text{AF}) \neq \emptyset$, i.e., there is at least one probability function that is both β -compliant and adheres to the set of rationality conditions (we address the case $\mathcal{P}_T(\text{AF}) \cap \mathcal{P}^\beta(\text{AF}) = \emptyset$ in the next section).

Define $\mathcal{P}_T^\beta(\text{AF}) = \mathcal{P}_T(\text{AF}) \cap \mathcal{P}^\beta(\text{AF})$.

Definition 1. Let β be a partial probability assignment and T a set of rationality conditions. Then the *possible probabilities* of $\mathcal{A} \in \text{Arg} \setminus \text{dom } \beta$, denoted as $\text{p}_{T, \text{AF}}^\beta(\mathcal{A})$, is defined as $\text{p}_{T, \text{AF}}^\beta(\mathcal{A}) = \{P(\mathcal{A}) \mid P \in \mathcal{P}_T^\beta(\text{AF})\}$.

Under the assumption $\mathcal{P}_T^\beta(\text{AF}) \neq \emptyset$, it is clear that $\text{p}_{T, \text{AF}}^\beta(\mathcal{A}) \neq \emptyset$ as well.

Example 3. We continue Example 2 with $\beta_1(\mathcal{S}_1) = 0.7$ and assume $T_1 = \{\text{COH}\}$. Then for the arguments $\mathcal{S}_2, \mathcal{I}, \mathcal{G}$ we obtain $\text{p}_{T_1, \text{AF}}^{\beta_1}(\mathcal{S}_2) = [0, 0.3]$, $\text{p}_{T_1, \text{AF}}^{\beta_1}(\mathcal{I}) = [0, 0.3]$, and $\text{p}_{T_1, \text{AF}}^{\beta_1}(\mathcal{G}) = [0, 0.7]$.

We need some set theoretical notions before we can state our next result. A subset X of a topological space is (*path*-)connected, if for every two elements $x, y \in X$ there is a continuous function $f : [0, 1] \rightarrow X$ with $f(0) = x$ and $f(1) = y$.² A set X is called *convex*, if for every two elements $x, y \in X$ and $\delta \in [0, 1]$ we also have $\delta x + (1 - \delta)y \in X$. A set X is *closed* if for every sequence $x_1, x_2, \dots \in X$, if $\lim_{n \rightarrow \infty} x_i$ exists then $\lim_{n \rightarrow \infty} x_i \in X$.

Proposition 2. Let AF be an abstract argumentation framework and β a partial probability assignment.

1. The set $\mathcal{P}^\beta(\text{AF})$ is connected, convex, and closed.
2. The sets $\mathcal{P}(\text{AF})$, $\mathcal{P}_{\text{COH}}(\text{AF})$, $\mathcal{P}_{\text{SFOU}}(\text{AF})$, $\mathcal{P}_{\text{FOU}}(\text{AF})$, $\mathcal{P}_{\text{OPT}}(\text{AF})$, $\mathcal{P}_{\text{SOPT}}(\text{AF})$, $\mathcal{P}_{\text{JUS}}(\text{AF})$ are connected, convex, and closed.

²Note that $\mathcal{P}(\text{AF})$ is a topological space as it can be identified with a subspace of $[0, 1]^n$ with $n = |\text{Arg}|$.

3. The set $\mathcal{P}_{\text{RAT}}(\text{AF})$ is connected and closed, but not convex in general.
4. For every $T \subseteq \{\text{COH}, \text{SFOU}, \text{FOU}, \text{OPT}, \text{SOPT}, \text{JUS}\}$ the sets $\mathcal{P}_T(\text{AF})$ and $\mathcal{P}_T^\beta(\text{AF})$ are connected, convex, and closed.
5. For every $T \subseteq \{\text{COH}, \text{SFOU}, \text{FOU}, \text{OPT}, \text{SOPT}, \text{JUS}\}$ and $\mathcal{A} \in \text{Arg}$ the set $\text{p}_{T, \text{AF}}^\beta(\mathcal{A})$ is connected, convex, and closed.

The final statement above is equivalent to saying that $\text{p}_{T, \text{AF}}^\beta(\mathcal{A})$ is an interval. Note also, that if $\text{RAT} \in T$ the set $\mathcal{P}_T(\text{AF})$ is closed, but not necessarily connected or convex. In the following, we focus on the cases where $T \subseteq \{\text{COH}, \text{SFOU}, \text{FOU}, \text{OPT}, \text{SOPT}, \text{JUS}\}$.

Proposition 2 implies that the problem of determining $\text{p}_{T, \text{AF}}^\beta(\mathcal{A})$ is equivalent to the classical *probabilistic entailment problem* (Jaumard, Hansen, and Poggi 1991; Lukasiewicz 2000). Due to space limitations, we do not go into more details here regarding this relationship but only exploit it to make some observations on the computational complexity of some problems related to $\text{p}_{T, \text{AF}}^\beta(\mathcal{A})$.

Proposition 3. Let AF be an abstract argumentation framework, β a partial probability assignment, and $T \subseteq \{\text{COH}, \text{SFOU}, \text{FOU}, \text{OPT}, \text{SOPT}, \text{JUS}\}$.

1. Deciding $p \in \text{p}_{T, \text{AF}}^\beta(\mathcal{A})$ for some $p \in [0, 1]$ is NP-complete.
2. Deciding $[l, u] = \text{p}_{T, \text{AF}}^\beta(\mathcal{A})$ for some $l, u \in [0, 1]$ is D^{P} -complete.
3. Computing $l, u \in [0, 1]$ such that $[l, u] = \text{p}_{T, \text{AF}}^\beta(\mathcal{A})$ is FP^{NP} -complete.

Besides using $\text{p}_{T, \text{AF}}^\beta(\mathcal{A})$ to obtain bounds on the probabilities of the remaining arguments, we might also be interested in obtaining point probabilities for the remaining arguments that are as unbiased as possible, giving the probabilistic information of β . One can use the *principle of maximum entropy* (see also Section 2.2) for this purpose, which is thanks to Proposition 2 also applicable in our context.

Definition 2. Let AF be an abstract argumentation framework, β a partial probability assignment, and $T \subseteq \{\text{COH}, \text{SFOU}, \text{FOU}, \text{OPT}, \text{SOPT}, \text{JUS}\}$. Define the set $\mathcal{P}_{\text{ME}}^{\beta, \text{AF}, T}$ via

$$\mathcal{P}_{\text{ME}}^{\beta, \text{AF}, T} = \arg \max_{Q \in \mathcal{P}_T^\beta(\text{AF})} H(Q)$$

Proposition 4. The set $\mathcal{P}_{\text{ME}}^{\beta, \text{AF}, T}$ contains exactly one uniquely defined probability function.

Due to the above proposition we identify the singleton set $\mathcal{P}_{\text{ME}}^{\beta, \text{AF}, T}$ with its only element, e.g., we write $\mathcal{P}_{\text{ME}}^{\beta, \text{AF}, T}(\mathcal{A})$ to denote the probability $P(\mathcal{A})$ with $\{P\} = \mathcal{P}_{\text{ME}}^{\beta, \text{AF}, T}$.

Example 4. We continue Examples 2 and 3 with $\beta_1(\mathcal{S}_1) = 0.7$ and assume $T_1 = \{\text{COH}\}$. Then we have $\mathcal{P}_{\text{ME}}^{\beta_1, \text{AF}, T_1}(\mathcal{S}_2) = 0.3$, $\mathcal{P}_{\text{ME}}^{\beta_1, \text{AF}, T_1}(\mathcal{I}) = 0.3$, and

$\mathcal{P}_{\text{ME}}^{\beta_1, \text{AF}, T_1}(\mathcal{G}) = 0.5$. Recall that $\mathcal{P}_{T_1, \text{AF}}^{\beta_1}(\mathcal{I}) = [0, 0.3]$ (Example 3) and observe that $\mathcal{P}_{\text{ME}}^{\beta_1, \text{AF}, T_1}(\mathcal{I}) = 0.3$ which is the maximal probability that can be assigned to \mathcal{I} . However, note also that this value is closest to 0.5 which is the probability value with the least amount of information (in the information-theoretic sense). Indeed, it can be observed that all probabilities assigned above are those closest to 0.5 which is a general feature of reasoning based on the principle of maximum entropy (note, however, that in more complex settings involving other rationality conditions the function $\mathcal{P}_{\text{ME}}^{\beta, \text{AF}, T}$ is not always characterized as easily as that).

Theorem 1 already pointed out the relationship of the principle of maximum entropy to grounded semantics. Note also that for β with $\text{dom } \beta = \emptyset$ and $T = \{\text{JUS}\}$ we have that $\mathcal{P}_{\text{ME}}^{\beta, \text{AF}, T}$ corresponds to the grounded labelling of AF. Taking into account partial probabilistic information we therefore extended the notion of a grounded labelling and obtain a probability function that is “as grounded as possible”. Similarly, if we exchange the maximum in Definition 2 by a minimum, we obtain a generalization of the notion of stable labelings, cf. Theorem 1 item 4.

4 Contradictory Probability Assessments

So far we assumed $\mathcal{P}_T^\beta(\text{AF}) \neq \emptyset$. In this section, we investigate the general scenario without this assumption. Consider the following example.

Example 5. We continue Example 2. New evidence obtained by analyzing the footage from camera \mathcal{S}_2 suggests that the probability of the person in those pictures is indeed John is 0.4. Therefore, the partial probability assessment β'_1 is defined by $\beta'_1(\mathcal{S}_1) = 0.7$ and $\beta'_1(\mathcal{S}_2) = 0.4$. Considering the set of rationality conditions $T = \{\text{COH}\}$ one can see that $\mathcal{P}_T^{\beta'_1}(\text{AF}) = \mathcal{P}_T(\text{AF}) \cap \mathcal{P}^{\beta'_1}(\text{AF}) = \emptyset$ as the condition

$$P(\mathcal{S}_1) \leq 1 - P(\mathcal{S}_2), \quad (1)$$

which is necessary for having $P \in \mathcal{P}_T(\text{AF})$, cannot be satisfied for any β'_1 -compliant P . In this case, one would still be interested in obtaining a “reasonable” probability for e. g. \mathcal{I} .

We address the issue outlined in the example above by, first, analyzing in a quantitative manner *how much* a partial probability assessment deviates from satisfying a given set of rationality conditions, and afterwards using this analysis to provide reasonable probabilities for the remaining arguments.

4.1 Inconsistency Measures for Contradictory Probability Assessments

We first address the question of how to measure the distance (or inconsistency) of a given partial probability assignment $\beta : \text{Arg} \rightarrow [0, 1]$ to the set $\mathcal{P}_T(\text{AF})$ of probability functions. As before we restrict our attention to $T \subseteq \{\text{COH}, \text{SFOU}, \text{FOU}, \text{OPT}, \text{SOPT}, \text{JUS}\}$. Let Π denote the set of all partial probability assignments and let \mathbb{A} denote the set of all abstract argumentation frameworks.

Definition 3. An *inconsistency measure* \mathcal{I}_T is a function $\mathcal{I}_T : \Pi \times \mathbb{A} \rightarrow [0, \infty)$.

The intuition behind an inconsistency measure \mathcal{I}_T is that for a partial probability assessment β and an argumentation framework AF, the value $\mathcal{I}_T(\beta, \text{AF})$ quantitatively assesses the severity of β violating the rationality conditions imposed by T in AF. In particular, larger values indicate greater violation while attaining the minimum $\mathcal{I}_T(\beta, \text{AF})$ suggests that β does not violate the rationality conditions imposed by T in AF at all. Note that inconsistency measures have been investigated before mostly in the context of classical logic, see e. g. (Hunter and Konieczny 2010).

Before formalizing the intuition behind an inconsistency measure, we need some further notation. Let β, β' be partial probability assignments. We say that β' is an *expansion* of β if $\text{dom } \beta \subseteq \text{dom } \beta'$ and $\beta(\mathcal{A}) = \beta'(\mathcal{A})$ for all $\mathcal{A} \in \text{dom } \beta$. If $\text{dom } \beta \cap \text{dom } \beta' = \emptyset$ then define $(\beta \circ \beta')$ with $\text{dom } \beta \circ \beta' = \text{dom } \beta \cup \text{dom } \beta'$ via $(\beta \circ \beta')(\mathcal{A}) = \beta(\mathcal{A})$ for $\mathcal{A} \in \text{dom } \beta$ and $(\beta \circ \beta')(\mathcal{A}) = \beta'(\mathcal{A})$ for $\mathcal{A} \in \text{dom } \beta'$. Two arguments $\mathcal{A}, \mathcal{B} \in \text{Arg}$ are (indirectly) connected if there is an undirected path between them. Let $CC(\text{AF})$ be the connected components of AF w. r. t. this notion of connectedness.

Now, some desirable properties for an inconsistency measure in our context—motivated by similar properties for inconsistency measures in classical logics (Hunter and Konieczny 2010)—are as follows.

Consistency $\mathcal{I}_T(\beta, \text{AF}) = 0$ iff $\mathcal{P}_T(\text{AF}) \cap \mathcal{P}^\beta(\text{AF}) \neq \emptyset$.

Monotonicity If β' is an expansion of β then $\mathcal{I}_T(\beta, \text{AF}) \leq \mathcal{I}_T(\beta', \text{AF})$.

Super-additivity If $\text{dom } \beta \cap \text{dom } \beta' = \emptyset$ then $\mathcal{I}_T(\beta \circ \beta', \text{AF}) \geq \mathcal{I}_T(\beta, \text{AF}) + \mathcal{I}_T(\beta', \text{AF})$.

Separability $\mathcal{I}_T(\beta, \text{AF}) = \sum_{\text{AF}' \in CC(\text{AF})} \mathcal{I}_T(\beta, \text{AF}')$.

The property *consistency* states that an inconsistency measure must attain its minimal value if and only if there is at least one β -compliant probability function P that satisfies all conditions T w. r. t. AF. The property *monotonicity* states that the inconsistency cannot decrease when adding further constraints to a partial probability assessment. The property *super-additivity* means that the sum of the inconsistency values of two independent probability assessments cannot be larger than the inconsistency value of the joint probability assessment. Finally, the property *separability* demands that the inconsistency value decomposes on the connected components of an argumentation framework.

In order to implement inconsistency measures for our setting of probabilistic argumentation, we base our measures on metrics on probability functions, cf. (Thimm 2013; De Bona and Finger 2015; Grant and Hunter 2013).

Definition 4. A function $d : \mathcal{P} \times \mathcal{P} \rightarrow [0, \infty)$ is called a *pre-metrical distance measure* if it satisfies $d(P, P') = 0$ if and only if $P = P'$.

In the following, we refer to *pre-metrical distance measures* simply by *distance measures* (note that we do not impose the properties *symmetry* and *triangle equality* of full

distance measures). Examples of such distance measures are (let $p \geq 1$)

$$d_{\text{KL}}(P, P') = \sum_{x \in \text{dom } P \cap \text{dom } P'} P(x) \log \frac{P(x)}{P'(x)}$$

$$d_p(P, P') = \sqrt[p]{\sum_{x \in \text{dom } P \cap \text{dom } P'} |P(x) - P'(x)|^p}$$

In the definition of d_{KL} , if $x = 0$ we assume $x \log x/y = 0$ and if $x \neq 0$ but $y = 0$ we assume $x \log x/y = x$. The measure d_{KL} is also called the *Kullback-Leibler divergence* (or *relative entropy*). The measure d_p is called the *p-norm distance*. In the following, we will use these two distance measures as examples to illustrate our approach. Note that any other pre-metrical distance measure can be used instead.

Note that both measures d_{KL} and d_p are defined over the set $\text{dom } P \cap \text{dom } P'$. This is only a technical necessity in order to have well-defined measures for all pairs of probability functions. In the following, distance measures are only applied on pairs of probability functions P and P' such that $P, P' \in \mathcal{P}(\text{AF})$ for some AF, i. e., $\text{dom } P = \text{dom } P'$.

For a distance measure d , a probability function $P \in \mathcal{P}$ and closed sets $\mathcal{Q}, \mathcal{Q}' \subseteq \mathcal{P}$ we abbreviate

$$d(P, \mathcal{Q}) = \min_{P' \in \mathcal{Q}} d(P, P')$$

$$d(\mathcal{Q}, P) = \min_{P' \in \mathcal{Q}} d(P', P)$$

$$d(\mathcal{Q}, \mathcal{Q}') = \min_{P' \in \mathcal{Q}'} d(P', \mathcal{Q})$$

Using a distance measure d on probability functions we define a general inconsistency measure as follows.

Definition 5. Let d be a distance measure and $T \subseteq \{\text{COH}, \text{SFOU}, \text{FOU}, \text{OPT}, \text{SOPT}, \text{JUS}\}$. The *distance-based* inconsistency measure $\mathcal{I}_T^d : \Pi \times \mathbb{A} \rightarrow [0, \infty)$ for T and d is defined via

$$\mathcal{I}_T^d(\beta, \text{AF}) = d(\mathcal{P}^\beta(\text{AF}), \mathcal{P}_T(\text{AF}))$$

In other words, $\mathcal{I}_T^d(\beta, \text{AF})$ is the minimal distance w. r. t. d of a β -compliant probability function P_1 and a probability function P_2 that satisfies the rationality conditions of T w. r. t. AF.

Example 6. We continue Example 5 with $T_1 = \{\text{COH}\}$ and β'_1 defined by $\beta'_1(\mathcal{S}_1) = 0.7$ and $\beta'_1(\mathcal{S}_2) = 0.4$. For d_1 (the Manhattan distance) it can be easily seen that $\mathcal{I}_{T_1}^{d_1}(\beta'_1, \text{AF}) = 0.1$ as this amounts to the absolute amount Equation (1) is violated. For d_2 (the Euclidean distance) we obtain $\mathcal{I}_{T_1}^{d_2}(\beta'_1, \text{AF}) \approx 0.037^3$. For d_{KL} (the Kullback-Leibler divergence) we obtain $\mathcal{I}_{T_1}^{d_{\text{KL}}}(\beta'_1, \text{AF}) \approx 0.625$. A geometrical interpretation for both d_2 and d_{KL} is hard to provide but compare those values to the values obtained for β_2 defined by $\beta_2(\mathcal{S}_1) = 0.8$ and $\beta_2(\mathcal{S}_2) = 0.9$: $\mathcal{I}_{T_1}^{d_1}(\beta_2, \text{AF}) = 0.7$, $\mathcal{I}_{T_1}^{d_2}(\beta_2, \text{AF}) \approx 0.403$, and $\mathcal{I}_{T_1}^{d_{\text{KL}}}(\beta_2, \text{AF}) \approx 0.312$. From

³Values of inconsistency measures were determined by using the OpenOpt optimization package <http://openopt.blogspot.de>

an intuitive point of view β_2 seems more inconsistent than β'_1 (the constraint (1) is violated to a larger extent) and both $\mathcal{I}_{T_1}^{d_1}(\beta_2, \text{AF})$ and $\mathcal{I}_{T_1}^{d_2}(\beta_2, \text{AF})$ comply with this intuition as they assign larger inconsistency values to β_2 than to β'_1 . For $\mathcal{I}_{T_1}^{d_{\text{KL}}}(\beta_2, \text{AF})$ we obtain the opposite result, due to the fact that d_{KL} does not measure distance of probabilities but distance of information content. However, we leave a deeper interpretation of this matter for future work.

As can be seen by the following results, the family of inconsistency measures \mathcal{I}_T^d complies with our formalization of a meaningful inconsistency measure.

Theorem 2. *If d is a pre-metrical distance measure then \mathcal{I}_T^d satisfies consistency.*

Theorem 3. *The function $\mathcal{I}_T^{d_{\text{KL}}}$ satisfies consistency and monotonicity.*

We conjecture that $\mathcal{I}_T^{d_{\text{KL}}}$ also satisfies *super-additivity* and *separability* (based on experiments), but a formal proof is yet to be found. For $\mathcal{I}_T^{d_p}$ we have a stronger result as follows.

Theorem 4. *For $p \geq 1$ the function $\mathcal{I}_T^{d_p}$ satisfies consistency and monotonicity. For $p = 1$ the function $\mathcal{I}_T^{d_p}$ also satisfies separability and super-additivity.*

For $p > 1$ a relaxed version of *separability* holds.

Theorem 5. *For $p > 1$ the function $\mathcal{I}_T^{d_p}$ satisfies*

$$\mathcal{I}_T^{d_p}(\beta, \text{AF}) \leq \sum_{\text{AF}' \in \text{CC}(\text{AF})} \mathcal{I}_T^{d_p}(\beta, \text{AF}')$$

4.2 Distance-based Consolidation

The measure \mathcal{I}_T^d allows to quantitatively assess the violation of a partial probability assignment in the light of a given set of rationality conditions. However, Example 5 suggests that even in the presence of contradictory information, we want to be able to provide reasonable inference results. Following the idea of \mathcal{I}_T^d we define the set of reasonable probability functions as those probability functions in $\mathcal{P}^\beta(\text{AF})$ that are *closest* to satisfying the rationality conditions T .

Definition 6. Define the set $\Pi_{T,d,\text{AF}}(\beta) \subseteq \mathcal{P}^\beta(\text{AF})$ via

$$\Pi_{T,d,\text{AF}}(\beta) = \{P \in \mathcal{P}^\beta(\text{AF}) \mid d(P, \mathcal{P}_T(\text{AF})) \text{ minimal}\}$$

In other words, the set $\Pi_{T,d,\text{AF}}(\beta)$ is the set of “witnesses” of the inconsistency value $\mathcal{I}_T^d(\beta, \text{AF})$, i. e., those probability functions P with $d(P, \mathcal{P}_T(\text{AF})) = \mathcal{I}_T^d(\beta, \text{AF})$.

Our idea is now to use $\Pi_{T,d,\text{AF}}(\beta)$ in the same way for reasoning as we used $\mathcal{P}_T^\beta(\text{AF})$ in Section 3. In fact, it can be easily seen that under the assumption $\mathcal{P}_T^\beta(\text{AF}) \neq \emptyset$ reasoning with $\Pi_{T,d,\text{AF}}(\beta)$ coincides with our previous approach.

Proposition 5. *If $\mathcal{P}_T^\beta(\text{AF}) \neq \emptyset$ then $\Pi_{T,d,\text{AF}}(\beta) = \mathcal{P}_T^\beta(\text{AF})$ for every pre-metrical distance measure d .*

Moreover, $\Pi_{T,d,\text{AF}}(\beta)$ is a strict generalization of $\mathcal{P}_T^\beta(\text{AF})$ as it always contains probability functions, even if $\mathcal{P}_T^\beta(\text{AF}) = \emptyset$. Furthermore, $\Pi_{T,d,\text{AF}}(\beta)$ features the same topological properties as $\mathcal{P}_T^\beta(\text{AF})$ if the distance measure d is reasonably chosen.

Proposition 6. $\Pi_{T,d,AF}(\beta) \neq \emptyset$.

Proposition 7. For strictly convex d and $T \subseteq \{\text{COH}, \text{SFOU}, \text{FOU}, \text{OPT}, \text{SOPT}, \text{JUS}\}$ the set $\Pi_{T,d,AF}(\beta)$ is connected, convex, and closed.

The above statement is true for our examples of distance measures, except for d_1 (the Manhattan distance) which is not strictly convex.

Corollary 1. For $T \subseteq \{\text{COH}, \text{SFOU}, \text{FOU}, \text{OPT}, \text{SOPT}, \text{JUS}\}$ and $d \in \{d_{KL}, d_p\}$ (for $p > 1$) we have that $\Pi_{t,d,AF}(\beta)$ is a connected, convex, and closed set.

The above results show that $\Pi_{T,d,AF}(\beta)$ behaves exactly like $\mathcal{P}_T^\beta(\text{AF})$ (in the topological sense) and is a strict generalization. We therefore extend the notion of *possible probabilities* $\mathcal{P}_{T,AF}^\beta(\mathcal{A})$ to the general case.

Definition 7. Let d be strictly convex and $T \subseteq \{\text{COH}, \text{SFOU}, \text{FOU}, \text{OPT}, \text{SOPT}, \text{JUS}\}$. Define

$$\pi_{T,AF}^{\beta,d}(\mathcal{A}) = \{P(\mathcal{A}) \mid P \in \Pi_{T,d,AF}(\beta)\}$$

Observe that by Proposition 5 we have $\pi_{T,AF}^{\beta,d}(\mathcal{A}) = \mathcal{P}_{T,AF}^\beta(\mathcal{A})$ if $\mathcal{P}_T^\beta(\text{AF}) \neq \emptyset$ (for every pre-metrical distance measure d).

Example 7. We continue Example 5 with $T_1 = \{\text{COH}\}$ and β'_1 defined by $\beta'_1(\mathcal{S}_1) = 0.7$ and $\beta'_1(\mathcal{S}_2) = 0.4$. For the Euclidean distance d_2 we obtain

$$\pi_{T_1,AF}^{\beta'_1,d_2}(\mathcal{I}) \approx [0.0284, 0.383] \quad \pi_{T_1,AF}^{\beta'_1,d_2}(\mathcal{G}) \approx [0.0270, 0.682]$$

which shows that beliefs in both \mathcal{I} and \mathcal{G} can be quite low (due to the conflict in the evidence) but that the belief in \mathcal{G} can be up to 0.682 due to the stronger evidence in \mathcal{S}_1 and weaker evidence in \mathcal{S}_2 .

While the problem of determining $\mathcal{P}_{T,AF}^\beta(\mathcal{A})$ is equivalent to the classical probabilistic entailment problem (see Section 3) the problem of determining $\pi_{T,AF}^{\beta,d}(\mathcal{A})$ has strong relationships with the generalized probabilistic entailment problem of (Potyka and Thimm 2015). In particular, the work of (Potyka and Thimm 2015) suggests that the computational complexity of determining $\pi_{T,AF}^{\beta,d}(\mathcal{A})$ (and the other related computational problems) is the same as for $\mathcal{P}_{T,AF}^\beta(\mathcal{A})$, provided that the inconsistency value $\mathcal{I}_T^d(\beta, \text{AF})$ has already been determined. We leave a deeper analysis of this matter for future work.

Similarly as for $\mathcal{P}_T^\beta(\text{AF})$ we can define reasoning based on maximum entropy on $\Pi_{T,d,AF}(\beta)$ as follows.

Definition 8. Let d be strictly convex and $T \subseteq \{\text{COH}, \text{SFOU}, \text{FOU}, \text{OPT}, \text{SOPT}, \text{JUS}\}$. Define

$$\mathcal{P}_{ME}^{\beta,AF,T,d} = \arg \max_{Q \in \Pi_{T,d,AF}(\beta)} H(Q)$$

The validity of the following proposition follows also straightforwardly from our previous results.

Proposition 8. Let AF be an abstract argumentation framework, β a partial probability assignment, d be strictly convex, and $T \subseteq \{\text{COH}, \text{SFOU}, \text{FOU}, \text{OPT}, \text{SOPT}, \text{JUS}\}$. The set $\mathcal{P}_{ME}^{\beta,AF,T,d}$ contains exactly one uniquely defined probability function.

We also write $\mathcal{P}_{ME}^{\beta,AF,T,d}(\mathcal{A})$ to denote the probability $P(\mathcal{A})$ with $\{P\} = \mathcal{P}_{ME}^{\beta,AF,T,d}$.

Example 8. We continue Example 7 with $T_1 = \{\text{COH}\}$ and β'_1 defined by $\beta'_1(\mathcal{S}_1) = 0.7$ and $\beta'_1(\mathcal{S}_2) = 0.4$. For the Euclidean distance d_2 we obtain

$$\mathcal{P}_{ME}^{\beta'_1,AF,T_1,d_2}(\mathcal{I}) \approx 0.3788 \quad \mathcal{P}_{ME}^{\beta'_1,AF,T_1,d_2}(\mathcal{G}) \approx 0.4959$$

which gives an ambiguous picture on the innocence or guilt of John (due to the contradictory information), with a slight tendency towards guilt due to the slightly higher belief in \mathcal{S}_1 .

4.3 An Alternative Point of View

In Definition 6 we defined $\Pi_{T,d,AF}(\beta)$ to be a subset of probability functions of $\mathcal{P}^\beta(\text{AF})$ that are closest to the set $\mathcal{P}_T(\text{AF})$. The decision of defining $\Pi_{T,d,AF}(\beta)$ like this was based on the need to consider only probability functions that are compliant with our observations but as rational as possible w. r. t. T . Consider now

$$\Pi_{T,d,AF}^*(\beta) = \{P \in \mathcal{P}_T(\text{AF}) \mid d(\mathcal{P}^\beta(\text{AF}), P) \text{ minimal}\}$$

The set $\Pi_{T,d,AF}^*(\beta)$ contains those probability functions in $\mathcal{P}_T(\text{AF})$ that are closest to the set $\mathcal{P}^\beta(\text{AF})$, i. e., probability functions that are fully rational w. r. t. T and maximally compliant with our observations. It can be easily seen that $\Pi_{T,d,AF}^*(\beta)$ behaves exactly like $\Pi_{T,d,AF}(\beta)$ w. r. t. its topological properties.

Proposition 9.

1. If $\mathcal{P}_T^\beta(\text{AF}) \neq \emptyset$ then $\Pi_{T,d,AF}^*(\beta) = \Pi_{T,d,AF}(\beta) = \mathcal{P}_T^\beta(\text{AF})$ for every pre-metrical distance measure d .
2. $\Pi_{T,d,AF}^*(\beta) \neq \emptyset$
3. For strictly convex d and $T \subseteq \{\text{COH}, \text{SFOU}, \text{FOU}, \text{OPT}, \text{SOPT}, \text{JUS}\}$ the set $\Pi_{T,d,AF}^*(\beta)$ is connected, convex, and closed.

Consequently, we can define reasoning based on $\Pi_{T,d,AF}^*(\beta)$ in the same way as we did on $\Pi_{T,d,AF}(\beta)$.

If we view the $\Pi_{T,d,AF}(\beta)$ and $\Pi_{T,d,AF}^*(\beta)$ operators as repair operators, then they offer us two options:

1. $\Pi_{T,d,AF}(\beta)$ is used when we want to preserve the prior information we have in β but want to get as close as possible to satisfying the rationality constraints in T ; and
2. $\Pi_{T,d,AF}^*(\beta)$ is used when we want to impose the rationality constraints we have in T but want to keep as much as possible from the prior information we have in β .

We can regard $\Pi_{T,d,AF}(\beta)$ as *soft repair* as it does not satisfy T but gets closer to it, and we can regard $\Pi_{T,d,AF}^*(\beta)$ as *hard repair* as it does satisfy T . So hard repairs ensure conformity with T but at the loss of some the original information in β , whereas soft repairs ensure no loss of the original information in β , but at the loss of some conformity with T .

5 Related Works

The two main approaches to probabilistic argumentation are the constellations and the epistemic approaches (Hunter 2013).

- In the constellations approach, the uncertainty is in the topology of the graph (see for example (Dung and Thang 2010; Li, Oren, and Norman 2011; Hunter 2012; Fazzinga, Flesca, and Parisi 2013; Li, Oren, and Norman 2013; Hunter and Thimm 2014a; Dondio 2014; Polberg and Doder 2014; Doder and Woltran 2014; Fazzinga, Flesca, and Parisi 2015)). As an example, this approach is useful when one agent is not sure what arguments and attacks another agent is aware of, and so this can be captured by a probability distribution over the space of possible argument graphs.
- In the epistemic approach, the topology of the argument graph is fixed, but there is uncertainty about whether an argument is believed (Thimm 2012; Hunter 2013; Baroni, Giacomin, and Vicig 2014; Hunter 2014b; 2014a; Hunter and Thimm 2014d; 2014c; 2014b; Hunter 2015). A core idea of the epistemic approach is that the more likely an agent is to believe in an argument, the less likely it is to believe in an argument attacking it.

This paper is a development of the epistemic approach with a focus on incomplete and inconsistent probability distributions. These problems were first raised in (Hunter 2013; Hunter and Thimm 2014d), but no systematic solutions to the problems were presented. In contrast in this paper, we have provided solutions based on well-justified notions of distance between probability distributions.

In quantifying disagreement between argument graphs, the distance between labellings has been considered in terms of the weighted sum of the number of labellings that differ (Booth et al. 2012). Various kinds of distance have also been considered in methods for epistemic enforcement in abstract argumentation (Baumann and Brewka 2010; Baumann 2012; Coste-Marquis et al. 2014), for revising argument graphs (Coste-Marquis, Konieczny, and Maily 2014a; 2014b), and for merging argument graphs (Coste-Marquis et al. 2007; Delobelle, Konieczny, and Vesic 2015). There are related proposals for belief revision in argumentation such as (Cayrol, de Saint-Cyr, and Lagasque-Schiex 2010; Gabbay and Rodrigues 2012; Bisquert et al. 2013; Diller et al. 2015) but they do not use distance measures.

Dung and Thang (Dung and Thang 2010) provided the first proposal to extend abstract argumentation with a probability distribution over sets of arguments which they use with a version of assumption-based argumentation in which a subset of the rules are probabilistic rules. Another approach to augmenting abstract argumentation with probabilities has used equations based on the structure of the graph to constrain the probability assignments, and these can be solved to calculate probabilities (Gabbay and Rodrigues 2015). In another rule-based system for argumentation, the belief in the premises of an argument is used to calculate the belief in the argument (Riveret et al. 2007). However, the proposal does not investigate further the nature of this

assignment, for example with respect to abstract argumentation, but rather its use in dialogue is explored. In a logic-based approach, Verheij combines qualitative reasoning in terms of reasons and defeaters, with quantitative reasoning using argument strength which is modeled as the conditional probability of the conclusions given the premises (Verheij 2014).

There are other approaches to bringing probability theory into systems for dialogical argumentation. A probabilistic model of the opponent has been used in a dialogue strategy allowing the selection of moves for an agent based on what it believes the other agent is aware of and the moves it might take (Rienstra 2012; Rienstra, Thimm, and Oren 2013). In another approach to probabilistic opponent modeling, the history of previous dialogues is used to predict the arguments that an opponent might put forward (Hadjinikolis et al. 2013). For modeling the possible dialogues that might be generated by a pair of agents, a probabilistic finite state machine can represent the possible moves that each agent can make in each state of the dialogue assuming a set of arguments that each agent is aware of (Hunter 2014b). This has been generalized to POMDPs when there is uncertainty about what an opponent is aware of (Hadoux et al. 2015).

Some research has investigated relationships between Bayesian networks and argumentation. Bayesian networks can be used to model argumentative reasoning with arguments and counterarguments (Vreeswijk 2004). In a similar vein, Bayesian networks can be used to capture aspects of argumentation in the Carneades model where the propagation of argument applicability and statement acceptability can be expressed through conditional probability tables (Grabmair, Gordon, and Walton 2010). Going the other way, arguments can be generated from a Bayesian network, and this can be used to explain the Bayesian network (Timmer et al. 2015), and argumentation can be used to combine multiple Bayesian networks (Nielsen and Parsons 2007).

6 Discussion and Summary

The epistemic approach provides a finer grained assessment of an argument graph than given by the basic notions of extensions. With labellings, arguments are labelled as in, out, or undec, whereas with the epistemic approach an argument can take any value in $[0, 1]$. By adopting constraints on the probability distribution, we have shown in previous work how the epistemic approach subsumes Dung's approach (Thimm 2012; Hunter and Thimm 2014c). However, we have also argued that there is a need for a view where we adopt weaker constraints on the probability distribution. For instance, an important aspect of the epistemic approach is the representation of disbelief in arguments even when they are unattacked. It is not always possible or practical to identify a counterargument to reject in argumentation, and often it is quite natural to directly represent the disbelief in an argument without consideration of the counterargument.

The epistemic approach is also useful for modeling the belief that an opponent might have in the arguments that could be presented, which is useful for example when deciding on the best arguments to present in order to persuade that opponent. Strategies in dialogical argumentation are an

important research issue (Thimm 2014). By harnessing a model of the beliefs of opponent, better choices can be made by an agent (see for example (Hunter 2015)).

In this paper, our focus has been on incomplete probability distributions and on probability distributions that are inconsistent with a set of constraints. These issues commonly arise when considering multiple agents. For instance, when using a probability distribution to represent the beliefs of an opponent, the opponent may have made explicit its beliefs in specific arguments (perhaps by positing them, or by answering queries regarding them). Normally, what is known about the beliefs of the opponent will be incomplete. To give an example of dealing with inconsistency, we can use the probability distribution to represent the feedback obtained from an audience of a television debate. Here, the probability distribution might be inconsistent with the chosen constraints. If we assume that the audience does conform to the constraints, and that probability distribution fails to satisfy the constraints, then we can “repair” the probability distribution, using our approaches of “soft” and “hard repair”.

References

- Baroni, P.; Giacomin, M.; and Vicig, P. 2014. On rationality conditions for epistemic probabilities in abstract argumentation. In *Proceedings of the Fifth International Conference on Computational Models of Argumentation (COMMA'14)*.
- Baumann, R., and Brewka, G. 2010. Expanding argumentation frameworks: Enforcing and monotonicity results. In *Proceedings of COMMA 2010*, 75–86.
- Baumann, R. 2012. What does it take to enforce an argument? minimal change in abstract argumentation. In *Proceedings of ECAI'12*, 127–132.
- Bisquert, P.; Cayrol, C.; de Saint-Cyr, F. D.; and Lagasquie-Schiex, M. 2013. Enforcement in argumentation is a kind of update. In *Scalable Uncertainty Management (SUM'13)*, volume 8078 of *LNCS*, 30–42. Springer.
- Booth, R.; Caminada, M.; Podlaszewski, M.; and Rahwan, I. 2012. Quantifying disagreement in argument-based reasoning. In Conitzer; Winikoff; Padgham; and van der Hoek., eds., *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*.
- Cayrol, C.; de Saint-Cyr, F. D.; and Lagasquie-Schiex, M.-C. 2010. Change in abstract argumentation frameworks: Adding an argument. *Journal of Artificial Intelligence Research* 38:49–84.
- Coste-Marquis, S.; Devred, C.; Konieczny, S.; Lagasquie-Schiex, M.; and Marquis, P. 2007. On the merging of dung’s argumentation systems. *Artificial Intelligence* 171(10-15):730–753.
- Coste-Marquis, S.; Konieczny, S.; Maily, J.-G.; and Marquis, P. 2014. Extension enforcement in abstract argumentation as an optimization problem. In *Proceedings of IJCAI'15*, 2876–2882.
- Coste-Marquis, S.; Konieczny, S.; and Maily, J.-G. 2014a. On the revision of argumentation systems: Minimal change of argument statuses. In *Proceedings of KR'14*, 72–81.
- Coste-Marquis, S.; Konieczny, S.; and Maily, J.-G. 2014b. A translation-based approach from revision of argumentation frameworks. In *Proceedings of JELIA'14*, 77–85.
- De Bona, G., and Finger, M. 2015. Measuring inconsistency in probabilistic logic: rationality postulates and dutch book interpretation. *Artificial Intelligence* 227:140–164.
- Delobelle, J.; Konieczny, S.; and Vesic, S. 2015. On the aggregation of argumentation frameworks. In *Proceedings of IJCAI'15*, 2911–2917.
- Diller, M.; Haret, A.; Linsbichler, T.; S.Rümmele; and Woltran, S. 2015. An extension-based approach to belief revision in abstract argumentation. In *Proceedings of IJCAI'15*, 2926–2932.
- Doder, D., and Woltran, S. 2014. Probabilistic argumentation frameworks – a logical approach. In Straccia, U., and Cali, A., eds., *Scalable Uncertainty Management*, volume 8720 of *Lecture Notes in Computer Science*. Springer International Publishing. 134–147.
- Dondio, P. 2014. Multi-valued and probabilistic argumentation frameworks. In *Proceedings of the Fifth International Conference on Computational Models of Argumentation (COMMA'14)*.
- Dung, P. M., and Thang, P. M. 2010. Towards (probabilistic) argumentation for jury-based dispute resolution. In *Proceedings of the 2010 conference on Computational Models of Argument: Proceedings of COMMA 2010*, 171–182. Amsterdam, The Netherlands, The Netherlands: IOS Press.
- Dung, P. M. 1995. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. *Artificial Intelligence* 77(2):321–358.
- Fazzinga, B.; Flesca, S.; and Parisi, F. 2013. On the complexity of probabilistic abstract argumentation. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI'13)*.
- Fazzinga, B.; Flesca, S.; and Parisi, F. 2015. On the complexity of probabilistic abstract argumentation frameworks. *ACM Transactions on Computational Logic* 16(3):22:1–22:39.
- Gabbay, D., and Rodrigues, O. 2012. A numerical approach to the merging of argumentation networks. In *Proceedings of CLIMA'12*, volume 7486 of *LNCA*, 195–212. Springer.
- Gabbay, D., and Rodrigues, O. 2015. Probabilistic argumentation: An equational approach. *Logica Universalis* 9(3):345–382.
- Grabmair, M.; Gordon, T.; and Walton, D. 2010. Probabilistic semantics for the carneades argument model using bayesian networks. In *Proceedings of COMMA*, 255–266.
- Grant, J., and Hunter, A. 2013. Distance-based measures of inconsistency. In *Proceedings of the 12th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'13)*, volume 7958 of *Lecture Notes in Computer Science*, 230–241. Springer.
- Hadjinikolis, C.; Siantos, Y.; Modgil, S.; Black, E.; and McBurney, P. 2013. Opponent modelling in persuasion dialogues. In *Proceedings of IJCAI*, 164170.

- Hadoux, E.; Beynier, A.; Maudet, N.; Weng, P.; and Hunter, A. 2015. Optimization of probabilistic argumentation with markov decision models. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI'15)*.
- Hunter, A., and Konieczny, S. 2010. On the measure of conflicts: Shapley inconsistency values. *Artificial Intelligence* 174(14):1007–1026.
- Hunter, A., and Thimm, M. 2014a. Probabilistic argument graphs for argumentation lotteries. In *Proceedings of the Fifth International Conference on Computational Models of Argumentation (COMMA'14)*.
- Hunter, A., and Thimm, M. 2014b. Probabilistic argumentation with epistemic extensions. In *Proceedings of the International Workshop on Defeasible and Ampliative Reasoning (DARE'14)*.
- Hunter, A., and Thimm, M. 2014c. Probabilistic argumentation with epistemic extensions and incomplete information. Technical report, ArXiv.
- Hunter, A., and Thimm, M. 2014d. Probabilistic argumentation with incomplete information. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI'14)*, 1033–1034.
- Hunter, A. 2012. Some foundations for probabilistic abstract argumentation. In *Computational Models of Argument (COMMA'12)*, 117–128. IOS Press.
- Hunter, A. 2013. A probabilistic approach to modelling uncertain logical arguments. *International Journal of Approximate Reasoning* 54(1):47–81.
- Hunter, A. 2014a. Probabilistic qualification of attack in abstract argumentation. *International Journal of Approximate Reasoning* 55:607–638.
- Hunter, A. 2014b. Probabilistic strategies in dialogical argumentation. In *Proceedings of the 8th International Conference on Scalable Uncertainty Management (SUM'14)*.
- Hunter, A. 2015. Modelling the persuadee in asymmetric argumentation dialogues for persuasion. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI'15)*.
- Jaumard, B.; Hansen, P.; and Poggi, M. 1991. Column generation methods for probabilistic logic. *ORSA - Journal on Computing* 3(2):135–148.
- Jaynes, E. T. 1957. Information theory and statistical mechanics. *Physical Review* 106:620–630.
- Kern-Isberner, G. 2001. Characterizing the principle of minimum cross-entropy. In *Conditionals in Nonmonotonic Reasoning and Belief Revision: Considering Conditionals as Agents, LNCS 2087*. Springer-Verlag Berlin Heidelberg, 73–90.
- Li, H.; Oren, N.; and Norman, T. J. 2011. Probabilistic argumentation frameworks. In *Proceedings of the First International Workshop on the Theory and Applications of Formal Argumentation (TFAFA'11)*.
- Li, H.; Oren, N.; and Norman, T. J. 2013. Relaxing independence assumptions in probabilistic argumentation. In *Proceedings of the Tenth International Workshop on Argumentation in Multi-Agent Systems (ArgMAS'13)*.
- Lukasiewicz, T. 2000. Probabilistic logic programming with conditional constraints. *ACM Transactions on Computational Logic* 2(3):289–339.
- Nielsen, S., and Parsons, S. 2007. An application of formal argumentation: Fusing bayesian networks in multi-agent systems. *Artificial Intelligence* 171(10-15):754–775.
- Paris, J. B. 1994. *The Uncertain Reasoner's Companion – A Mathematical Perspective*. Cambridge University Press.
- Polberg, S., and Doder, D. 2014. Probabilistic abstract dialectical frameworks. In *Proceedings of the 14th European Conference on Logics in Artificial Intelligence (JELIA'14)*.
- Potyka, N., and Thimm, M. 2015. Probabilistic reasoning with inconsistent beliefs using inconsistency measures. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI'15)*.
- Rienstra, T.; Thimm, M.; and Oren, N. 2013. Opponent models with uncertainty for strategic argumentation. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI'13)*.
- Rienstra, T. 2012. Towards a probabilistic dung-style argumentation system. In *Proceedings AT2012*.
- Riveret, R.; Rotolo, A.; Sartor, G.; Prakken, H.; and Roth, B. 2007. Success chances in argument games: a probabilistic approach to legal disputes. In *Legal Knowledge and Information Systems (JURIX07)*, 99108. IOS Press.
- Thimm, M. 2012. A probabilistic semantics for abstract argumentation. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI'12)*.
- Thimm, M. 2013. Inconsistency measures for probabilistic logics. *Artificial Intelligence* 197:1–24.
- Thimm, M. 2014. Strategic argumentation in multi-agent systems. *Künstliche Intelligenz, Special Issue on Multi-Agent Decision Making* 28(3):159–168.
- Timmer, S. T.; Meyer, J.-J. C.; Prakken, H.; Renooij, S.; and Verheij, B. 2015. Explaining bayesian networks using argumentation. In *Proceedings of the 13th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'15)*.
- Verheij, B. 2014. Arguments and their strength: Revisiting pollock's anti-probabilistic starting points. In *Proceedings of the Fifth International Conference on Computational Models of Argumentation (COMMA'14)*.
- Vreeswijk, G. 2004. Argumentation in bayesian belief networks. In *Proceedings of ArgMAS*, 111–129.
- Wu, Y., and Caminada, M. 2010. A labelling-based justification status of arguments. *Studies in Logic* 3(4):12–29.