Artificial Intelligence Research and Development
M. Villaret et al. (Eds.)

doi:10.3233/FAIA210111

On Probabilistic Logical Argumentation Based on Conditional Probability

Pilar DELLUNDE [a], Lluís GODO [b] and Amanda VIDAL [b]
[a] *Universitat Autònoma de Barcelona (UAB) and Barcelona Graduate School of Mathematics (BGSMath), Bellaterra, Spain*
[b] *Artificial Intelligence Research Institute (IIIA-CSIC), Bellaterra, Spain*

Abstract. In this paper, we introduce a framework for probabilistic logic-based argumentation inspired on the DeLP formalism and an extensive use of conditional probability. We define probabilistic arguments built from possibly inconsistent probabilistic knowledge bases and study the notions of attack, defeat and preference between these arguments. Finally, we discuss consistency properties of admissible extensions of the Dung's abstract argumentation graphs obtained from sets of probabilistic arguments and the attack relations between them.

Keywords. Probabilistic Argumentation; Logic-based Argumentation; Defeasible Knowledge; Inconsistent Probabilistic Knowledge Bases.

1. Introduction

In many scenarios, one has to deal with both uncertain and inconsistent information. Argumentation systems [4] have shown to be very suitable tools to reason with inconsistent information. In the literature, there have been a number of approaches to combine different theories of argumentation, both abstract and instantiated, with probability theory and other uncertainty models in order to allow for a more fine-grained reasoning when arguments involve uncertain information. See for instance [1,9,15,21,11,12,13,2,14,19,6].

According to [11], two main approaches can be distinguished. The epistemic approach takes the stance that the uncertainty is within the instantiated arguments (probabilities are used for capturing the strength of an argument, given the uncertainty concerning the truth of its premises or the reliability of its inferences, see for instance [20]). In the constellation approach, usually in the frame of abstract argumentation, the uncertainty is about the arguments themselves (probabilities are used for expressing uncertainty about the acceptance of the argument by some arguing agent, see for instance [11,12]). In contrast to [11], but similarly to [19], in this paper we consider logic-based arguments $A = (\mathit{support}; \mathit{conclusion})$, where *support* and *conclusion* are logical propositions, pervaded with uncertainty due a non-conclusive conditional link between their supports and their conclusions. In such a case, it is very reasonable to supplement the argument representation with a quantification α of how certain *conclusion* can be claimed to hold whenever *support* is known to hold [17], leading to represent an arguments as triples $A = (\mathit{support}; \mathit{conclusion} : \alpha)$. A very natural choice is to interpret α as some parameter related to the conditional probability $P(\mathit{conclusion} \mid \mathit{support})$. In this paper we will consider α to be a probability interval $[c_1, c_2]$, meaning that the argument A provides the information $P(conclusion \mid support) \in [c_1, c_2]$.

If we internalise the conditional link within the argument as a conditional formula *support* $\rightsquigarrow$ *conclusion* and arguments get more complex and need several uncertain conditionals to link the support with the conclusion, then we can attach conditional probability intervals to each of the involved conditionals, so arguments become of the form

$$A = (\Pi, \Delta = \{(\psi_1 \rightsquigarrow \varphi_1 : \beta_1), \ldots, (\psi_n \rightsquigarrow \varphi_n : \beta_n)\};\ \varphi\colon\ \alpha),$$

where Π is a finite set of factual (i.e. non conditional) premises, ψ_i, φ_i's are logical propositions, β_i's are probability intervals, and α is a probability interval with which φ can be logically entailed from Π and Δ. In fact, this type of arguments can be seen as a probabilistic generalization of those at work in the Defeasible Logic Programming argumentation framework (DeLP) [10].

In this paper we extend our preliminary work in [7]. First of all, we consider a more general language to build arguments, allowing conditionals (or rules) to have arbitrary propositional formulas as antecedents and consequents, while in [7] only conjunction of literals and literals were allowed. Second, we allow to attach intervals to conditionals for the corresponding conditional probabilities rather than only lower bounds as in [7]. Finally, in this paper we start the study of Dung's abstract argumentation systems associated to our probabilistic argumentation systems, and in particular the status of Caminada and Amgoud's rationality postulates [5] for the acceptability semantics based on complete extensions.

This paper is structured as follows. In Section 2 we introduce the notions about logic and probability necessary for the rest of the paper. In Section 3 we present our framework of probabilistic argumentation based on conditional probabilities, with the notion of probabilistic arguments and the attack and defeat relations among them. In Section 5 we consider the abstract argumentation system associated to a set of probabilistic arguments and their attacks, and we study the status of rationality postulates for Dung's extensions-based complete semantics. We conclude the paper with some comments on future work.

2. Logic and probability

When aiming towards the definition of a formal argumentation framework, a first step is the selection of the underlying language and logical system that will govern the derivation of new knowledge from a given set of information. Let $\mathcal{L}$ be a classical propositional language built over a finite set of variables $\mathcal{V}$ and $\vdash$ be the consequence relation of classical propositional logic.

Following [16], we introduce the set of probabilistic conditionals. We let $\mathcal{L}_{P_r}$ be the set of expressions of the form $\psi \rightsquigarrow \phi$: $[c_1, c_2]$, where ψ, ϕ are formulas of $\mathcal{L}$ and $c_1 \leq c_2$ are real numbers from the unit interval $[0, 1]$. We call ϕ the *consequent* of $\psi \rightsquigarrow \phi$: $[c_1, c_2]$, and ψ its *antecedent*. We can distinguish between classical and purely probabilistic conditionals. *Classical* conditionals are either of the form $\psi \rightsquigarrow \phi$: $[1, 1]$ or of the form $\psi \rightsquigarrow \phi$: $[0, 0]$, and purely probabilistic conditionals are of the form $\psi \rightsquigarrow \phi$: $[c_1, c_2]$, with $c_1 < 1$ and $c_2 > 0$.

Let Ω stand for the (finite) set of classical truth-evaluations $e : \mathcal{L} \to \{0, 1\}$ of the formulas in $\mathcal{L}$. Probabilities on the set of formulas $\mathcal{L}$ can be introduced in the standard way, as it is done in probability logics [6], namely by defining a probability distribution on the set of interpretations Ω, and extending it to all formulas by adding up the probabilities of their models. In other words, given a probability distribution $\pi : \Omega \to [0,1]$, i.e. is such that $\sum_{e\in\Omega} \pi(e) = 1$, then π induces a probability on formulas $P : \mathcal{L} \to [0,1]$ by stipulating

$$P(\varphi) = \sum_{e\in\Omega : e(\varphi)=1} \pi(e).$$

As defined, P satisfies the classical axioms of probability measures, namely $P(\top) = 1$ and finite additivity $P(\phi \vee \psi) = P(\phi) + P(\psi)$ whenever $\phi \wedge \psi \vdash \bot$, and moreover P respects logical equivalence: if $\vdash \phi \leftrightarrow \psi$ then $P(\phi) = P(\psi)$. Conversely, for any probability P on $\mathcal{L}$, there is a probability distribution π on Ω such that π induces P as above. Given a probabilistic conditional $\psi \rightsquigarrow \phi$: $[c_1, c_2] \in \mathcal{L}_{P_r}$, we say that a probability P on $\mathcal{L}$ *satisfies* $\psi \rightsquigarrow \phi$: $[c_1, c_2]$, if either $P(\psi) = 0$, or $P(\psi) > 0$ and $P(\phi \mid \psi) := P(\phi \wedge \psi)/P(\psi) \in [c_1, c_2]$, written

$$P \models_{pr} \psi \rightsquigarrow \phi\text{: } [c_1, c_2],$$

namely, the conditional probability of ϕ given ψ is a number in the real interval $[c_1, c_2]$. Remark that the probability of a conditional $\psi \rightsquigarrow \varphi$ is interpreted as the conditional probability $P(\phi \mid \psi)$, not as the probability of the material implication $P(\neg\psi \vee \phi)$. Moreover, we say that P *satisfies* a set of probabilistic conditionals Σ, if it satisfies each expression in Σ. We will denote the set of probabilities that satisfy Σ by $PMod(\Sigma)$.

Now we introduce a consequence relation on the set of probabilistic conditionals $\mathcal{L}_{P_r}$ first defined in [16].

Definition 2.1. *For any subset of probabilistic conditionals* $\Sigma \cup \{\psi \rightsquigarrow \phi\text{: } [c_1, c_2]\} \subseteq \mathcal{L}_{P_r}$, *the conditional* $\psi \rightsquigarrow \phi$: $[c_1, c_2]$ *is a consequence of* Σ, *denoted by* $\Sigma \models_{pr} \psi \rightsquigarrow \phi$: $[c_1, c_2]$, *if every probability* $P \in PMod(\Sigma)$ *satisfies* $\psi \rightsquigarrow \phi$: $[c_1, c_2]$.

Notice that, so defined, $\models_{pr}$ on $\mathcal{L}_{P_r}$ is a Tarskian consequence relation, satisfying the rules of reflexivity, monotonicity and cut.

3. Using conditional probability in arguments

3.1. Knowledge bases and arguments

Usually, arguments are built from an initial knowledge base from which one can build arguments pro and against certain pieces of information non explicitly contained in the knowledge base. Our notion of knowledge base is inspired by the approach of DeLP and other argumentation systems oriented to work with not fully reliable information. The encoding of the knowledge about a given domain in these systems distinguish pieces of knowledge considered as certain and consistent (strict knowledge) and knowledge that is tentative and subject to uncertainty or inconsistency (defeasible knowledge). If probabilities are added to the pieces of uncertain knowledge, a finer separation arises, contributing to the trustworthiness and accurateness of arguments and their relations, and allowing the argumentation system to produce more detailed outputs.

The intuition is that strict knowledge about a domain is always consistent and certain information, and hence it can be implicitly used in any argument. Thus, to specify an argument, it is only needed to specify which observations or factual information are assumed, and which part of the uncertain probabilistic knowledge is based upon. In this work, we assume the strict domain knowledge to be consistent knowledge with probability equal to 1.

Definition 3.1. *For any set of propositional formulas $S \subseteq \mathcal{L}$, let the closure of S under a set of probabilistic conditionals $\Sigma \subseteq \mathcal{L}_{P_r}$ (denoted by $Cl_\Sigma(S)$) be the smallest set containing S and the consequent of any rule in Σ whose antecedent φ is such that $Cl_\Sigma(S) \vdash \varphi$, where $\vdash$ is the consequence relation of classical propositional logic.*

Inspired in Def. 2 from [18] we distinguish two different notions of consistency.

Definition 3.2 (Direct and indirect consistency)**.** *A set $S \subseteq \mathcal{L}$ is* directly consistent *iff $S \nvdash \bot$, and* indirectly consistent *with respect to a set Γ of classical probabilistic conditionals if $Cl_\Gamma(S)$, the closure of S under Γ, is directly consistent.*

Following the usual terminology in the field of argumentation, we call *strict rule* a classical expression in $\mathcal{L}_{P_r}$ (*facts* are strict rules of the form $\top \rightsquigarrow \phi$: $[1,1]$, that we will simply denote sometimes as ϕ:$[1,1]$) and *probabilistic defeasible rule* a purely probabilistic expression in $\mathcal{L}_{P_r}$. If Π is a set of strict rules, we will denote by Π_f the set of facts in Π and we will let $\Pi_r = \Pi \setminus \Pi_f$ be the set of proper strict rules. Moreover, we will also let $\overline{\Pi}_f = \{\phi \mid \top \rightsquigarrow \phi\text{: } [1,1] \in \Pi_f\} \subset \mathcal{L}$ the set of consequents of the facts.

Definition 3.3. *A* probabilistic knowledge base *is a pair $KB = (\Pi, \Delta)$, where Π is a finite set of consistent strict rules and Δ is a finite set of probabilistic defeasible rules.*

Example 3.4. *The following set of probabilistic (strict and defeasible) rules $\Pi \cup \Delta$ encodes generic knowledge about inferring whether a damage (a big monetary loss) is caused in a house when the evidence (in Π) is that the alarm goes off, possibly because of a burglary or a fire.*

$$\Pi = \begin{Bmatrix} alarm\colon\ [1,1] \\ fire \rightsquigarrow mon\_loss\colon\ [1,1] \end{Bmatrix}, \quad \Delta = \begin{Bmatrix} alarm \rightsquigarrow burglary\colon\ [0.6, 0.8] \\ alarm \rightsquigarrow fire\colon\ [0.4, 0.5] \\ burglary \rightsquigarrow mon\_loss\colon\ [0.9, 1] \\ burglary \wedge alarm \rightsquigarrow \neg mon\_loss\colon\ [0.8, 1] \end{Bmatrix}$$

With the tools introduced above, we propose the following notion for a probabilistic logic-based argument.

Definition 3.5 (Argument)**.** *Given a probabilistic knowledge base $KB = (\Pi, \Delta)$, a* probabilistic argument *$\mathcal{A}$ for a formula $\theta \in \mathcal{L}$, is a tuple $\mathcal{A} = (\Sigma, \theta, [c_1, c_2])$, where $\Sigma \subseteq \Delta$, and such that:*

- *Probabilistic argument consistency: $PMod(\Sigma \cup \Pi) \neq \emptyset$*
- *Logical adequacy: θ belongs to the closure of the set of consequents of Π_f under the rules of $\Sigma \cup \Pi_r$, i.e. $\theta \in Cl_{\Sigma \cup \Pi_r}(\overline{\Pi}_f)$.*

- c_1 *(resp.* c_2*) is the infimum (resp. the supremum) of the values* $P(\theta)$ *with* $P \in PMod(\Pi \cup \Sigma)$. *In other words,*
  $c_1 = \sup\{d_1 \in [0,1] \mid \Pi \cup \Sigma \models_{Pr} \top \rightsquigarrow \theta\colon [d_1, 1]\}$,
  $c_2 = \inf\{d_2 \in [0,1] \mid \Pi \cup \Sigma \models_{Pr} \top \rightsquigarrow \theta\colon [0, d_2]\}$.
- Σ *is minimal satisfying the above conditions.*

If $c_2 > 0.5$ *we will say the argument is* proper.

In the above definition we have chosen to model evidences in an argument as part of the strict knowledge, and thus as literals with probability 1, rather than events on which to compute the conditional probability of the argument conclusion θ, see e.g. [7] where both options are considered. Of course, this issue is debatable and we let a discussion for future work. Given an argument $\mathcal{A} = (\Sigma, \theta, [c_1, c_2])$ w.r.t. $KB = (\Pi, \Delta)$, we will denote by $\Pi_{\mathcal{A}}$ the set of minimal subsets of Π such that $\mathcal{A}$ is still an argument w.r.t. $KB^* = (\Pi', \Delta)$ for each $\Pi' \in \Pi_{\mathcal{A}}$. That is, $\Pi_{\mathcal{A}}$ gathers all the minimal sets of strict rules of Π really needed in the argument $\mathcal{A}$.[1]

Thus, an argument for a literal provides for both a logical and an optimal probabilistic derivation of its conclusion from its premises. In this we follow [11] in decoupling the logic and the probabilistic aspects. Note that, in a sense, the requirements of the existence of a logical derivation and of a probabilistic entailment are rather independent. For instance, if p, q, r are variables, let $\Sigma = \{p \rightsquigarrow q\colon [\alpha, 1], q \rightsquigarrow r\colon [\beta, 1]\}$. If $0 < \alpha, \beta < 1$, then we have $r \in Cn_{\Sigma}(p)$, but $\{\top \rightsquigarrow p\colon [1,1]\} \cup \Sigma \not\models_{Pr} \top \rightsquigarrow r\colon [\gamma, 1]$, for any $\gamma > \alpha \cdot \beta$. Conversely, $\{\top \rightsquigarrow \neg q\colon [1,1]\} \cup \Sigma \models_{Pr} \top \rightsquigarrow \neg p\colon [1,1]$, but $\neg p \notin Cn_{\Sigma}(\neg q)$.

Some examples of probabilistic arguments over the KB from Example 3.4 are:

$$\mathcal{A}_1 = (\{alarm \rightsquigarrow burglary\colon [0.6, 0.8], burglary \rightsquigarrow mon\_loss\colon [0.9, 1]\}; mon\_loss\colon [0.54, 1])$$

$$\mathcal{A}_2 = (\{alarm \rightsquigarrow burglary\colon [0.6, 0.8], burglary \wedge alarm \rightsquigarrow \neg mon\_loss\colon [0.8, 1]\}; \neg mon\_loss\colon [0.48, 1])$$

$$\mathcal{A}_3 = (\{alarm \rightsquigarrow fire\colon [0.4, 0.5]\}; mon\_loss\colon [0.4, 0.5])$$

Note that $\mathcal{A}_1$ and $\mathcal{A}_2$ are proper, while $\mathcal{A}_3$ is not.

3.2. Counter-arguments, attacks and defeats

The next step in our formalization of a probabilistic argumentation system is to introduce the notions of subargument and attack relation between arguments. The notion of subargument is the usual one.

Definition 3.6 (Subargument)**.** *Let* $\mathcal{A} = (\Sigma, \theta, [c_1, c_2])$ *be an argument. A subargument of* $\mathcal{A}$ *is an argument* $\mathcal{B} = (\Sigma', \theta', [d_1, d_2])$ *where* $\Sigma' \subseteq \Sigma$.

Next we identify when two probabilistic conditionals lead to an inconsistency in the context of a KB.

Definition 3.7 (Disagreement)**.** *Let* $KB = (\Pi, \Delta)$ *be a probabilistic knowledge base. We say that two conditionals* $\psi_1 \rightsquigarrow \varphi_1\colon [c_1, c_2]$ *and* $\psi_2 \rightsquigarrow \varphi_2\colon [d_1, d_2]$ disagree *whenever they are probabilistically inconsistent with the strict knowledge, i.e. when* $PMod(\Pi \cup \{\psi_1 \rightsquigarrow \varphi_1\colon [c_1, c_2], \psi_2 \rightsquigarrow \varphi_2\colon [d_1, d_2]\} = \emptyset$.

[1] It can be shown that the set $\Pi_{\mathcal{A}}$ is not necessarily a singleton.

This notion is used to define the following attack relation between arguments that is somehow more general than an undercut relation (for a general reference on this type of relations we refer to [4]).

Definition 3.8 (Attack). *An argument* $\mathcal{A} = (\Sigma_1, \theta_1, [c_1, c_2])$ *attacks another argument* $\mathcal{B} = (\Sigma_2, \theta_2, [d_1, d_2])$ *at a formula* α *if there is a subargument* $\mathcal{B}' = (\Sigma_2', \alpha, [e_1, e_2])$ *of* $\mathcal{B}$ *such that* $\top \rightsquigarrow \theta_1 \colon [c_1, c_2]$ *and* $\top \rightsquigarrow \alpha \colon [e_1, e_2]$ *disagree.*

Going back to the examples above, it is clear that the argument $\mathcal{A}_2$ attacks both arguments $\mathcal{A}_1$ and $\mathcal{A}_3$, but it turns out that arguments $\mathcal{A}_1$ and $\mathcal{A}_3$ also attack each other despite of they both conclude on the same formula, the reason being that $\top \rightsquigarrow mon\_loss \colon [0.9, 1]$ and $\top \rightsquigarrow mon\_loss \colon [0.4, 0.5]$ are probabilistically inconsistent.

The next question is to specify when an attack can be deemed as effective, that is, when an attack of argument $\mathcal{A}$ to another argument $\mathcal{B}$ invalidates the latter, or in other words when an attack is actually a defeat. In our case, having probabilities in the arguments provides an additional criterion with an important role to be played. One possibility is to directly use the involved weights to decide when an argument prevails over another one. For instance, according to this criterion, argument $\mathcal{A}_1$ defeats argument $\mathcal{A}_2$. However, this seems rather counter-intuitive since, even if the derived probability in $\mathcal{A}_2$ is smaller than the one derived in $\mathcal{A}_1$, argument $\mathcal{A}_2$ is using more information than $\mathcal{A}_1$. This is in essence the specificity criterion, an approach that has been developed in non probabilistic scenarios, e.g. in [10] or [3]. Nevertheless, in case of two conflicting arguments using the same amount of information, then the comparison of the probability values surely becomes a suitable criterion to use. We start by formalizing the notion of specificity following ideas of [10].

Definition 3.9 (Activation set). *Given a probabilistic knowledge base* $KB = (\Pi, \Delta)$*, an* activation set *of an argument* $\mathcal{A} = (\Sigma, \theta, [c_1, c_2])$ *is a set of formulas of the form* $Act = Ant(\Sigma \cup \Pi')$*, where* $Ant(\Sigma \cup \Pi')$ *is the set of the antecedents of all the rules in* Σ *and* Π'*, for some* $\Pi' \in \Pi_{\mathcal{A}}$*. We will denote by* $Act(\mathcal{A})$ *the set of all activation sets for* $\mathcal{A}$*.*

For instance, continuing with Example 3.4, we have that

$$
\begin{aligned}
Act(\mathcal{A}_1) &= \{\{alarm\}, \{burglary\}\},\\
Act(\mathcal{A}_2) &= \{\{alarm\}, \{burglary \wedge alarm\}\},\\
Act(\mathcal{A}_3) &= \{\{alarm\}, \{fire\}\}.
\end{aligned}
$$

In the following, if Γ and Γ' are two sets of formulas, we will write $\Gamma \vdash \Gamma'$ when $\Gamma \vdash \psi$ for all $\psi \in \Gamma'$. Further, if S and S' are two sets of sets of formulas, we will write $S \vdash S'$ when for all $\Gamma \in S$ there is $\Gamma' \in S'$ such that $S \vdash \Gamma'$.

Definition 3.10 (Specificity). *We say that an argument* $\mathcal{A}$ *is* more specific than *another argument* $\mathcal{B}$ *when* $Act(\mathcal{A}) \vdash Act(\mathcal{B})$ *but* $Act(\mathcal{B}) \nvdash Act(\mathcal{A})$*.* $\mathcal{A}$ *and* $\mathcal{B}$ *are* equi-specific *if both* $Act(\mathcal{A}) \vdash Act(\mathcal{B})$ *and* $Act(\mathcal{B}) \vdash Act(\mathcal{A})$*, and* incomparable *whenever* $Act(\mathcal{A}) \nvdash Act(\mathcal{B})$ *and* $Act(\mathcal{B}) \nvdash Act(\mathcal{A})$*.*

It is clear that according to the above definition, concerning the arguments following Example 3.4, $\mathcal{A}_2$ is more specific than $\mathcal{A}_1$, but both are incomparable with $\mathcal{A}_3$.

In order to compare two arguments, we will then give the specificity criterion the highest priority, and if it does not produce a proper comparison, the degrees of probability will determine the strongest argument. This decision is based on the fact that probabilities of the defeasible rules inside each argument, and of the argument itself, reflect the faithfulness of each rule and indirectly, of the argument itself.

Definition 3.11 (Strength)**.** *An argument* $\mathcal{A} = (\Sigma_1, \theta_1, [c_1, c_2])$ *is* stronger *than another argument* $\mathcal{B} = (\Sigma_2, \theta_2, [d_1, d_2])$ *when* $\mathcal{A}$ *is more specific than* $\mathcal{B}$*, or otherwise when* $\mathcal{A}$ *and* $\mathcal{B}$ *are equi-specific or incomparable and* $c_1 > d_1$.

From the previous notions of the attack and comparative strength, it is now natural to formalize when indeed an argument defeats or rebuts another argument.

Definition 3.12 (Defeat)**.** *An argument* $\mathcal{A} = (\Sigma_1, \theta_1, [c_1, c_2])$ defeats *another argument* $\mathcal{B} = (\Sigma_2, \theta_2, [d_1, d_2])$ *when* $\mathcal{A}$ *attacks* $\mathcal{B}$ *on a subargument* $\mathcal{B} = (\Sigma_2', \alpha, [e_1, e_2])$ *and* $\mathcal{A}$ *is stronger than* $\mathcal{B}'$.

Continuing with Example 3.4, we have that $\mathcal{A}_2$ defeats $\mathcal{A}_1$ because $\mathcal{A}_2$ is more specific than $\mathcal{A}_1$, and $\mathcal{A}_2$ defeats $\mathcal{A}_3$ because $\mathcal{A}_2$ is not comparable with $\mathcal{A}_3$ and its probability lower bound is greater than that of $\mathcal{A}_3$. On the other hand $\mathcal{A}_1$ defeats $\mathcal{A}_3$ too, since they are also non-comparable and the probability lower bound of $\mathcal{A}_1$ is greater than the probability lower bound of $\mathcal{A}_3$.

4. Probabilistic Abstract Argumentation Frameworks

In this section we study Dung's abstract argumentation systems associated to our probabilistic argumentation systems. Namely, to any probabilistic conditional knowledge base $KB = (\Pi, \Delta)$ and set $\mathfrak{P}$ of *proper* probabilistic arguments over KB,[2] we can associate an abstract argumentation system $\langle \mathfrak{P}, R \rangle$, where R is the attack relation of Def. 3.8 on the set $\mathfrak{P}$. This is an argument graph in the sense of Dung's abstract argumentation [8].

In order to decide whether to accept a set of arguments, one can consider whether this set is conflict-free, admissible, or more in general if it fits with any of extension-based semantics proposed in the literature (complete, grounded, preferred, stable, etc). In [5] the authors define three principles, called *rationality postulates*, that can be used to identify some important properties of closure and consistency of argumentation systems. In this section, we prove that, in general, only one of these postulates is satisfied in our system. In the final section we propose some ideas in order to improve our system in a way that can satisfy the other two postulates.

Following [5], we restrict our language assuming that in the expressions of the form $\psi \rightsquigarrow \phi$: $[c_1, c_2]$, ψ can be only a literal, and ϕ a conjunction of literals. We recall now some basic notions of abstract argumentation theory. First we introduce the notion of extension under Dung's standard semantics, and then the notion of conflict-free and complete extension.

Definition 4.1 (Extension)**.** *Given an abstract argumentation framework* $\langle \mathfrak{P}, R \rangle$*, an* extension *of* $\langle \mathfrak{P}, R \rangle$ *is a subset of arguments* $E \subseteq \mathfrak{P}$*. Moreover, it is said that an extension* E defends *an argument* $\mathcal{A}$*, if for every argument* $\mathcal{B}$ *that attacks* $\mathcal{A}$*, there is an argument* $C \in E$ *that attacks* $\mathcal{B}$*.*

[2] Such a set $\mathfrak{P}$ somehow corresponds to what Hunter calls a *epistemic* extension in [11].

Given an argument $\mathcal{A}$ we denote by $Conc(\mathcal{A})$ the conclusion of $\mathcal{A}$. For instance, in case of a probabilistic argument $\mathcal{A} = (\Sigma, l, [c_1, c_2])$, $Conc(\mathcal{A}) = l$ is a literal. For any extension E, we denote by $Concs(E)$ the set of conclusions of the arguments in E.

Definition 4.2 (Conflict-free and complete extension)**.** *Given an abstract argumentation framework $\langle \mathfrak{P}, R\rangle$, it is said that an extension of $\langle \mathfrak{P}, R\rangle$, $E \subseteq \mathfrak{P}$ is* conflict-free *if there are not arguments $\mathcal{A}, \mathcal{B} \in E$ such that $\mathcal{A}R\mathcal{B}$, that is, such that $\mathcal{A}$ attacks $\mathcal{B}$. Moreover, it is said that E is* complete *if it is conflict-free and E defends all its arguments.*

Now we introduce the rationality postulates of [5]. The idea is to define postulates not only for each individual extension, but also for the set of overall justified conclusions, that is, the set called *Output* in [5], defined as follows: given the set of all extensions $E_1, \ldots, E_n$ under a given abstract semantics, $Output = \bigcap_{i \leq n} Concs(E_i)$.

In what follows $\mathfrak{P}$ will denote a set of probabilistic arguments over a probabilistic conditional knowledge base $KB = (\Pi, \Delta)$, and $\langle \mathfrak{P}, R\rangle$ will stand for the associated abstract argumentation system, as described at the beginning of this section.

Definition 4.3 (Cf. [5])**.** *Let $E_1, \ldots, E_n$ be the set of all extensions of $\langle \mathfrak{P}, R\rangle$ under a given abstract semantics. Then we have the following definitions:*

Postulate 1: $\langle \mathfrak{P}, R\rangle$ satisfies Closure *if*

1. $Concs(E_i) = Cl_\Pi(Concs(E_i))$*, for each $i \leq n$.*
2. $Output = Cl_\Pi(Output)$*.*

Postulate 2: $\langle \mathfrak{P}, R\rangle$ satisfies Direct Consistency *iff*

1. $Concs(E_i)$ *is consistent, for each $i \leq n$.*
2. $Output$ *is consistent.*

Postulate 3: $\langle \mathfrak{P}, R\rangle$ satisfies Indirect Consistency *iff*

1. $Cl_\Pi(Concs(E_i))$ *is consistent, for each $i \leq n$.*
2. $Cl_\Pi(Output)$ *is consistent.*

Next we show that our system satisfies Postulate 2, but not in general Postulates 1 and 3. Before we prove some basic properties regarding closure under subarguments and direct consistency of complete extensions.

Proposition 4.4. *For every complete extension E of $\langle \mathfrak{P}, R\rangle$, any argument $\mathcal{A} \in E$ and subargument $\mathcal{A}'$ of $\mathcal{A}$, we have that $\mathcal{A}' \in E$.*

Proof. Assume, searching for a contradiction, that $\mathcal{A} \in E$, $\mathcal{A}'$ is a subargument of $\mathcal{A}'$ but $\mathcal{A}' \notin E$. Then, since E is complete, either $E \cup \{\mathcal{A}`\}$ is not conflict-free, or $E$ does not defend $\mathcal{A}'$. In the first case, if $E \cup \{\mathcal{A}`\}$ is not conflict-free, there is a $\mathcal{B} \in E$ such that either $\mathcal{B}$ attacks $\mathcal{A}'$, or $\mathcal{A}'$ attacks $\mathcal{B}$. If $\mathcal{B}$ attacks $\mathcal{A}'$, by Def. 3.8, $\mathcal{B}$ attacks also $\mathcal{A}$, contradicting the fact that E is conflict-free. If $\mathcal{A}'$ attacks $\mathcal{B}$, since E is complete, E defends $\mathcal{B}$, and thus, there is a $C \in E$ that attacks $\mathcal{A}'$. Therefore, by Def. 3.8, C attacks also $\mathcal{A}$, contradicting the fact that E is conflict-free. In the second case, if E does not defend $\mathcal{A}'$, there is a $\mathcal{D} \in \mathfrak{P}$ that attacks $\mathcal{A}'$ and no $C \in E$ attacks $\mathcal{D}$. Then, by Def. 3.8, $\mathcal{D}$ attacks also $\mathcal{A}$ and E does not defend $\mathcal{A}$, which is a contradiction, because E is complete.

Since in both cases we have reached a contradiction, we can conclude that $\mathcal{A}' \in E$, i.e., E is closed under subarguments. □

Proposition 4.5. *For every conflict-free extension E of $\langle \mathfrak{P}, R\rangle$, $Concs(E)$ is directly consistent.*

Proof. Assume, searching for a contradiction, that $Concs(E)$ is not directly consistent, i.e. by Def. 3.2, $Concs(E) \vdash \bot$. Since $Concs(E)$ is a set of literals, there are arguments $\mathcal{A} = (\Sigma_1, l_1, [c_1, c_2])$, and $\mathcal{B} = (\Sigma_2, l_2, [d_1, d_2])$, with $\mathcal{A}, \mathcal{B} \in E$ and such that $l_1 = -l_2$.

Since E is conflict-free, $\mathcal{A}$ does not attack $\mathcal{B}$, and $\mathcal{B}$ does not attack $\mathcal{A}$. Thus, by Def. 3.8, $PMod(\Pi \cup \{(l_1 \mid \top)[c_1, c_2], (l_2 \mid \top)[d_1, d_2]\}) \neq \emptyset$, with $c_1, d_1 > 0.5$. Then $l_1 \neq -l_2$, contradicting our original assumption. Therefore $Concs(E)$ is directly consistent. □

Theorem 4.6. *If $\{E_1, \ldots, E_n\}$ is the set of all the extensions of $\langle \mathfrak{P}, R\rangle$ under the complete semantics, then $\langle \mathfrak{P}, R\rangle$ satisfies Postulate 2 but not necessarily Postulates 1 and 3.*

Proof. (Postulate 2) By Prop. 4.5, because for each $i \leq n$, E_i is complete, and thus conflict-free. Moreover, since the intersection of consistent sets is also consistent, we have that $Output$ is also consistent.

(Postulates 1 and 3) Let $KB = (\Pi, \Delta\rangle)$ be the following probabilistic conditional knowledge base:

$$\Pi = \{\top \rightsquigarrow d\colon [1,1], \top \rightsquigarrow e\colon [1,1], \top \rightsquigarrow f\colon [1,1], a \wedge b \rightsquigarrow \neg c\colon [1,1]\}$$

$$\Delta = \{r_1 : d \rightsquigarrow a\colon [0.9,1], r_2 : e \rightsquigarrow b\colon [0.7,1], r_3 : f \rightsquigarrow c\colon [0.8,1]\}$$

Consider now the following set $\mathfrak{P}$ of proper arguments over KB:

$\mathcal{A}_1 = (\{r_1\}, a, [0.9, 1])$, $\mathcal{A}_2 = (\{r_2\}, b, [0.7, 1])$,
$\mathcal{A}_3 = (\{r_3\}, c, [0.8, 1])$, $\mathcal{A}_4 = (\{r_1, r_2\}, \neg c, [0.6, 1])$.

On this set of arguments, the relation of attack is the following: $\mathcal{A}_3$ attacks $\mathcal{A}_4$, and $\mathcal{A}_4$ attacks $\mathcal{A}_3$. Thus, the abstract argumentation system associated to this set of arguments has the two following complete extensions: $E = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3\}$ and $F = \{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_4\}$. This is a counterexample for both Postulates 1 and 3, because E is not closed under strict rules and it is not indirectly consistent. □

5. Future work

In this paper we have proposed a framework for logic-based probabilistic argumentation where conditional expressions are qualified with conditional probability intervals, building on Lukasiewicz's setting for probabilistic logic programming with conditional constraints [16], and extending previous work [7].

There are many open problems left for future work. Here we mention some of them. First of all, although the underlying language based on conditionals is quite general, we could consider a more powerful logic to reason with those conditionals. We also need to consider and study possible alternatives to the notion of subargument, for instance in the line of [2] where a more refined notion is at work, and to the attack and defeat relations with possibly more suitable comparison criteria. Also observe that interval-valued probabilistic conditionals can be equivalently expressed by pairs of two lower bound-valued conditionals, which would be an interesting research line. Referring to this latter issue, and its influence in the fulfilment of rationality postulates by the associated abstract argumentation systems studied in the previous section, a promising prospect seems to consider a notion of collective conflict, also similarly to [2]. Finally, we can also mention the question of studying whether the probabilities involved in the arguments could allow for gradual notions of attack and acceptability.

Acknowledgments The authors are grateful to the reviewers and acknowledge partial support by the H2020-MSCA-RISE-2020 project MOSAIC (Grant Agreement 101007627) and the Spanish project ISINC (PID2019-111544GB-C21).

References

[1] T. Alsinet, C.I. Chesñevar, L. Godo, and G.R. Simari. A logic programming framework for possibilistic argumentation: Formalization and logical properties. *Fuzzy Sets and Systems*, 159(10):1208-1228, 2008.

[2] T. Alsinet, R. Béjar, L. Godo, F. Guitart. RP-DeLP: a weighted defeasible argumentation framework based on a recursive semantics. *J. Log. Comput.* 26(4): 1315-1360, 2016.

[3] D. Bamber, I.R. Goodman, and H.T. Nguyen. Robust reasoning with rules that have exceptions: From second-order probability to argumentation via upper envelopes of probability and possibility plus directed graphs. *Ann Math Artif Intell.*, 45:83–171, 2005.

[4] P. Baroni, D.M. Gabbay, M. Giacomin, and L. van der Torre (eds.) Handbook of Formal Argumentation. College Publications, 2018.

[5] M. Caminada and L. Amgoud. On the evaluation of argumentation formalisms. *Artificial Intelligence*, 171:286–310, 2007.

[6] F. Cerutti and M. Thimm. A general approach to reasoning with probabilities. *International Journal of Approximate Reasoning*, 111:35–50, 2019.

[7] P. Dellunde, L. Godo, and A. Vidal. Probabilistic Argumentation: An Approach Based on Conditional Probability -A Preliminary Report. In Proc. JELIA'21, *LNAI*, vol 12678, 25-32, 2021.

[8] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming, and n–person games. *Artificial Intelligence*, 77(2):321–357, 1995.

[9] P. M. Dung and P. M. Thang. Towards probabilistic argumentation for jury-based dispute resolution. In P. Baroni et al. (eds.), *Proc. of COMMA 2010*, vol. 216 of FAIA, pages 171–182. IOS Press, 2010.

[10] A. Garcia and G. Simari. Defeasible logic programming: an argumentative approach. *Theory and Practice of Logic Programming*, 4(1-2):95–138, 01 2004.

[11] A. Hunter. A probabilistic approach to modelling uncertain logical arguments. *International Journal of Approximate Reasoning*, 54(1):47–81, 2013.

[12] A. Hunter. Probabilistic qualification of attack in abstract argumentation. *IJAR*, 55(2):607–638, 2014.

[13] A. Hunter and M. Thimm. On partial information and contradictions in probabilistic abstract argumentation. In C. Baral et al. (eds.), *Proc. of KR 2016*, pages 53–62. AAAI Press, 2016.

[14] A. Hunter and M. Thimm. Probabilistic reasoning with abstract argumentation frameworks. *J. Artif. Intell. Res.*, 59:565–611, 2017.

[15] H. Li, N. Oren, and T. J. Norman. Probabilistic argumentation frameworks. In S. Modgil et al. (eds.), *Theory and Applications of Formal Argumentation - First International Workshop, TAFA 2011, Revised Selected Papers*, Vol. 7132 of *LNCS*, pages 1–16. Springer, 2011.

[16] T. Lukasiewicz. Probabilistic logic programming with conditional constraints. ACM Transactions on Computational Logic, 2(3): 289-339 (2001)

[17] J.L. Pollock. Justification and defeat. *Artificial Intelligence*, 67:377–408, 1994.

[18] H. Prakken. Historical overview of formal argumentation. In P. Baroni et al. (eds.) *Handbook of Formal Argumentation*, volume 1, pages 73–141. College Publications, 2018.

[19] H. Prakken. Probabilistic strength of arguments with structure. In M. Thielscher et al. (eds.), Francesca Toni, and Frank Wolter, editors, Proc. of KR 2018, pages 158–167. AAAI Press, 2018.

[20] S. Timmer, J.J.Ch. Meyer, H. Prakken, S. Renooij, and B. Verheij. A two-phase method for extracting explanatory arguments from bayesian networks. *Int. J. Approx. Reason.*, 80:475–494, 2017.

[21] B. Verheij. Jumping to conclusions: a logico-probabilistic foundation for defeasible rule-based arguments. In A. Herzig et al. (eds.), *Proc. of JELIA 2012*, LNAI, vol. 7519, 411–423, 2012.