ECAI 2020
G.D. Giacomo et al. (Eds.)

doi:10.3233/FAIA200165

Generating Instantiated Argument Graphs from Probabilistic Information

Anthony Hunter [1]

Abstract. The epistemic approach to probabilistic argumentation assigns belief to arguments. To better understand this approach, we consider structured arguments. Our approach is to start with a probability distribution, and generate an argument graph containing structured arguments with a probability assignment. We construct arguments directly from the probability distribution, rather than a knowledgebase, and then consider methods for selecting the arguments and counterarguments to present in the argument graph. This provides mechanisms for managing uncertainty in argumentation, and for argument-based explanations of probability distributions (that might come from data or from beliefs of an agent).

1 INTRODUCTION

Recently there has been some interest in extending abstract argumentation to handle probabilistic formalisation of uncertainty. Probabilistic approaches for modeling uncertainty in argumentation include the constellations approach and the epistemic approach [12]. The first is based on a probability distribution over the subgraphs of the argument graph ([11] which extends [5] and [19]), and this can be used to represent the uncertainty over the structure of the graph (i.e. whether a particular argument or attack appears in the argument graph under consideration). The second approach is the epistemic approach which involves a probability distribution over the subsets of the arguments [30, 12, 17]. This can be used to represent the uncertainty over which arguments are believed to be accepted. A further approach is based on labellings for arguments using *in*, *out*, and *undecided*, from [3], augmented with *off* for denoting that the argument does not occur in the graph [28]. A probability distribution over labellings gives a form of probabilistic argumentation that overlaps with the constellations and epistemic approaches.

The epistemic approach can be constrained (using axioms or postulates) to be consistent with Dung's dialectical semantics, but it can also be used as a potentially valuable alternative to Dung's dialectical semantics [30, 12, 17]. The later is advantageous if we want to consider enthymemes (arguments with incomplete support or claim) which create uncertainty about the exact meaning of the argument and associated attacks, and which in turns creates uncertainty about which arguments are acceptable. The epistemic approach also has advantages if we want to model imperfect agents who might not always adhere to the strict constraints of Dung' dialectical semantics as shown in [22]. However, focusing on the abstract level leaves open various questions as to the meaning and role of probabilities. For an epistemic argument, what is the meaning of the probability assignment? How can we instantiate a probabilistic argument with structure so that we can better understand its meaning? How can we generate logical arguments from probabilistic information? We are interested in these kinds of question in this paper.

For an argument A, $P(A)$ represents the degree of belief that A is acceptable. This belief is a function of the belief in the composition of the argument itself (for example, as a function of the belief in its premises, belief in its claim, and belief in the derivation of the claim from the premises), and the belief in the acceptability of other relevant arguments (e.g. directly/indirectly supporting or attacking). How we might formalize this depends on the kinds of arguments we are dealing with and the kinds of application. We will investigate one approach to formalizing this in this paper.

In previous work on probabilistic versions of structured argumentation [10, 12, 26], it is assumed that a knowledgebase of formulae is available, that the arguments are constructed from this knowledgebase, and that the probability values are then assigned to the arguments. A single probability distribution is assumed over the atoms of the language, and this can be used to give an assignment to arguments in such a way as to adhere some basic postulates for probabilistic argumentation. However, this still leaves open questions of how we get the knowledgebase, and what is a principled way of constructing the argument graph from the available arguments.

In order to address these questions in this paper, we take a different approach that uses the probability distribution to construct the structured arguments. This means that the probability distribution is primary, and the logical arguments and the argument graph constructed from them are secondary. This will give us a form of logical argumentation for the epistemic approach to probabilistic argumentation. We start with a probability distribution that reflects either the beliefs of the agent or the frequency distribution over some domain. From this probabilistic information, we investigate ways to construct structured argument graphs. First we consider an exhaustive approach which provides an unnecessarily complex argument graph, and then we consider ways to select arguments for inclusion in an argument graph. To do this in a principled way, we also consider desirable properties of the structured argument graphs that are constructed by selecting arguments. This provides the following contributions: (1) a probabilistic semantics for logical argumentation; (2) insights into selectivity in argumentation; (3) a framework for handling uncertainty in argumentation, in particular how the probabilities can be harnessed, and as such provides a step towards the learning of argument graphs from data; and (4) a formalism for explaining probability distributions (such as obtained by machine learning).

2 FRAMEWORK FOR ARGUMENTATION

We now propose a new form of logical argumentation called **essential argumentation** which is based on a language of literals. We assume a set of **literals** $\mathcal{L}$ formed from a set of **atomic propositions**

[1] University College London, UK, email: anthony.hunter@ucl.ac.uk

(i.e. **atoms**) $\mathcal{S}$. Each literal is either an atomic proposition α or its negation $\neg\alpha$. For an atom α, let $\mathsf{Atoms}(\alpha) = \mathsf{Atoms}(\neg\alpha) = \alpha$.

We will construct directed graphs (as proposed by Dung [6]). For a graph G, $\mathsf{Nodes}(G)$ is the set of nodes in G, and $\mathsf{Arcs}(G)$ is the set of arcs in G. Each node in a graph denotes an argument, and each arc (B, A) denotes an attack by B on A. An argument A is a **source** iff there is no argument B such that B attacks A. For graphs G, G', G is a **subgraph** of G', denoted $G \sqsubseteq G'$, iff $\mathsf{Nodes}(G) \subseteq \mathsf{Nodes}(G')$ and $\mathsf{Arcs}(G) \subseteq \{(A, B) \in \mathsf{Arcs}(G') \mid A, B \in \mathsf{Nodes}(G)\}$.

2.1 Probability distributions

We assume that a **probability distribution** P is over subsets of $\mathcal{S}$ (i.e. $P : \wp(\mathcal{S}) \rightarrow [0, 1]$ such that $\sum_{X \subseteq \mathcal{L}} P(X) = 1$). We use each subset of atoms as a model, and so for $X \subseteq \mathcal{S}$, and a Boolean combination of literals ϕ, we use $X \models \phi$ to denote X satisfies ϕ using the classical entailment relation. For example, $\{b\} \models \neg a$ where $\mathcal{S} = \{a, b\}$. For each $\mathcal{S}$, we assume an ordering over the sequence of atoms $\langle \alpha_1, \ldots, \alpha_n \rangle$ so that we can encode each model by a binary number: For a model X, if the ith argument is in X, then the ith digit is 1, otherwise it is 0. For example, for the sequence of atoms $\langle a, b, c \rangle$, the model $\{a, c\}$ is represented by 101.

The probability distribution concerns the truth of the atomic propositions. For example, if we are considering a domain with the concepts birds (b) and flying things (f), then we have a probability distribution over the powerset of $\{b, f\}$, and then we need to give a probability to each of $\{\}$, $\{b\}$, $\{f\}$, and $\{b, f\}$ being true with the sum being 1. We can do this by assuming a frame of reference. For instance, for birds flying, ways of determining the probability distribution include: (1) consider all **situations** (or observations, or days) to determine which situations have birds, birds flying, etc; (2) consider all **individuals** that are birds (or all the birds you have seen, or read about, or watched on TV) to determine which are flying, etc; or (3) consider all **types** of birds to determine which have the capability to fly (rather than you see actually fly) to determine which types have the capability to fly.

2.2 Essential arguments and attacks

We start with the simplest definition for an argument based on a set of literals where the support of an argument is a set of literals, and the claim is a literal.

Definition 1. *Let $\mathcal{L}$ be a set of literals. A* **naive argument** *is a tuple $\langle \Phi, \alpha \rangle$ where $\Phi \subseteq \mathcal{L}$ and $\alpha \in \mathcal{L}$. Let $\mathsf{Naive}(\mathcal{L})$ be the set of naive arguments formed from $\mathcal{L}$.*

In this paper, we regard a naive argument $\langle \{\beta_1, \ldots, \beta_n\}, \alpha \rangle$, as saying that an instance of "$\beta_1, \ldots, \beta_n$ is an instance of a". For example, for $A = \langle \{b, p\}, f \rangle$, where b denotes *bird*, p denotes *penguin*, and f denotes *flying-thing* we say that an instance of a bird and penguin is an instance of flying thing. A naive argument $\langle \Phi, \alpha \rangle$ is a **presumption** iff $\Phi = \emptyset$. For example, for $\langle \{\}, f \rangle$, any instance is a flying thing by presumption.

Definition 2. *For a pair of arguments $A_1 = \langle \Phi_1, \alpha_1 \rangle$ and $A_2 = \langle \Phi_2, \alpha_2 \rangle$, A_1* **attacks** *A_2 iff A_1 undercuts A_2 or A_1 rebuts A_2 where A_1* **undercuts** *A_2 iff $\beta \in \Phi_2$ such that $\{\alpha_1\} \models \neg\beta$, and A_1* **rebuts** *A_2 iff $\{\alpha_1\} \models \neg\alpha_2$.*

We could consider further types of counterarguments adapted from deductive argumentation [8]). The following graph definition presents all possible arguments and attacks without consideration of a probability distribution.

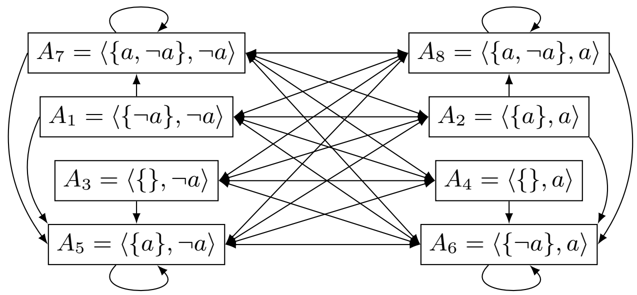

Figure 1: The super-exhaustive graph for $\mathcal{S} = \{a\}$.

Definition 3. *A* **super-exhaustive graph** *for a set of atoms $\mathcal{S}$ is a graph G where $\mathsf{Nodes}(G) = \mathsf{Naive}(\mathcal{L})$ and $\mathsf{Arcs}(G) = \{(A, B) \in \mathsf{Nodes}(G) \times \mathsf{Nodes}(G) \mid A \text{ attacks } B\}$.*

An example of a super-exhaustive graph is Figure 1. As seen here, any super-exhaustive graph contains inconsistent arguments (i.e. contradiction between support and claim) and reflexive arguments (i.e. claim is a premise). Hence, the graph contains a large number of arguments as quantified below.

Proposition 1. *For a language with n atoms (i.e. $|\mathcal{S}| = n$), the number of naive arguments is $2n \times 2^{2n}$.*

As defined next, essential arguments are naive arguments that are consistent and not reflexive. Returning to Figure 1, only A_3 and A_4 are essential arguments.

Definition 4. *Let $\mathcal{L}$ be a set of literals and let $\vdash$ be a classical consequence relation. An* **essential argument** *is a tuple $\langle \Phi, \alpha \rangle$ where $\Phi \subseteq \mathcal{L}$ and $\alpha \in \mathcal{L}$ such that (1) $\Phi \not\vdash \bot$ and (2) $\mathsf{Atoms}(\Phi) \cap \mathsf{Atoms}(\alpha) = \emptyset$. Let the set of essential arguments be $\mathsf{Args}(\mathcal{L})$. For $A = \langle \Phi, \alpha \rangle$, $\mathsf{Support}(A) = \Phi$, $\mathsf{Claim}(A) = \alpha$, and $\mathsf{Atoms}(A) = \mathsf{Atoms}(\Phi) \cup \mathsf{Atoms}(\alpha)$.*

In this paper, we use literals for the support and claim. However, for presenting them, we can give logically equivalent formulae for the support with claim. For instance, we could represent each support as a set of literals together with a formula that implies the claim. In other words for an essential argument of the form $\langle \{\beta_1, \ldots, \beta_n\}, \alpha \rangle$, we can give the logical argument $\langle \{\beta_1, \ldots, \beta_n, \beta_1 \wedge \ldots \wedge \beta_n \rightarrow \alpha\}, \alpha \rangle$. Clearly, the support of this is logically equivalent to the support plus claim $\{\beta_1, \ldots, \beta_n, \alpha\}$ of the essential argument.

Example 1. *Consider the essential arguments $A_1 = \langle \{b\}, f \rangle$ and $A_2 = \langle \{p, b\}, \neg f \rangle$. So A_1 (resp. A_2) is logically equivalent to $A'_1 = \langle \{b, b \rightarrow f\}, f \rangle$ (resp. $A_2 = \langle \{p, b, b \wedge p \rightarrow \neg f\}, \neg f \rangle$).*

Definition 5. *An* **exhaustive graph** *for a set of atoms $\mathcal{S}$ is a graph G where $\mathsf{Nodes}(G) = \mathsf{Args}(\mathcal{L})$ and $\mathsf{Arcs}(G) = \{(A, B) \in \mathsf{Args}(G) \times \mathsf{Args}(G) \mid A \text{ attacks } B\}$.*

We give an example of an exhaustive graph in Figure 2 for a language with two atoms. Because we have eliminated the inconsistent and reflexive arguments, we have the following result which is a relative improvement on the number of naive arguments, though still overwhelming when presented in an argument graph.

Proposition 2. *For a language with n atoms (i.e. $|\mathcal{S}| = n$), the number of essential arguments is $2n \times 3^{n-1}$.*

To address the problem of the exhaustive graph being so large, we introduce selectivity in Section 3. This will draw on the probability of the arguments as considered in the next subsection.

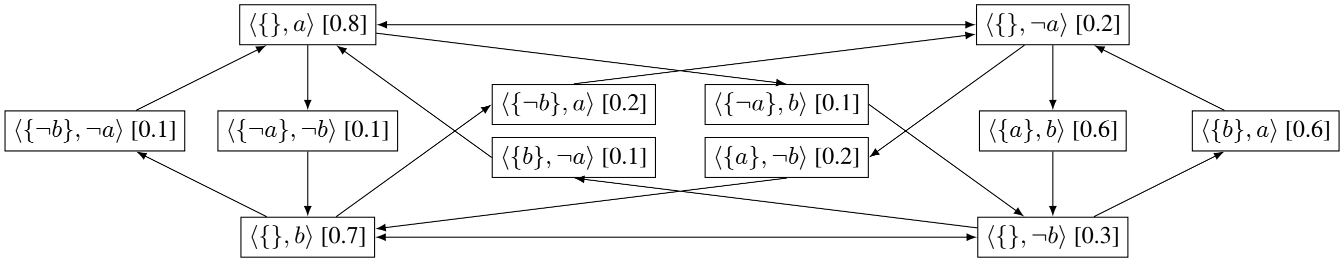

Figure 2: The exhaustive graph for the set of atoms $\mathcal{S} = \{a, b\}$. For $\langle a, b\rangle$, let $P(11) = 0.6$, $P(10) = 0.2$, $P(01) = 0.2$, and $P(00) = 0$. From this, the probability of each argument is given in square brackets using Definition 6.

2.3 Probability of an argument

The probability of an essential argument being acceptable is the belief in the premises and the claim. The same definition is used for deductive arguments in [12].

Definition 6. *The* **probability of an essential argument** $\langle\{\beta_1, \ldots, \beta_n\}, \alpha\rangle$, *denoted* $P(\langle\{\beta_1, \ldots, \beta_n\}, \alpha\rangle)$, *is* $\sum_{X \subseteq \mathcal{S} s.t. X \models \beta_1 \wedge \ldots \wedge \beta_n \wedge \alpha} P(X)$.

Example 2. *For the probability distribution, where b denotes* "bird"*, p denotes* "penguin"*, and f denotes* "flying-thing"*,* $A_1 = \langle\{b\}, f\rangle$ *has* $P(A_1) = 0.95$ *and* $A_2 = \langle\{p, b\}, \neg f\rangle$ *has* $P(A_2) = 0.01$.

$\langle b, p, f\rangle$	*110*	*101*	*100*
p	*0.01*	*0.95*	*0.04*

The probability $P(A_1) = 0.95$ *means that the belief in the argument being acceptable is very high.*

Proposition 3. *For arguments* $A, B \in \mathsf{Args}(\mathcal{L})$*, and probability distribution* P*, if* $\mathsf{Support}(A) \subseteq \mathsf{Support}(B)$*, and* $\mathsf{Claim}(A) = \mathsf{Claim}(B)$*, then* $P(A) \geq P(B)$.

We now return to the naive arguments that are inconsistent. As we show next, they have zero probability, and so gives a further reason to ignore them.

Proposition 4. *Let* P *be a probability distribution. For* $\langle\Phi, \alpha\rangle \in \mathsf{Naive}(\mathcal{L})$*, if* $\Phi \cup \{\alpha\} \vdash \bot$*, then* $P(\langle\Phi, \alpha\rangle) = 0$*, and* $\langle\Phi, \alpha\rangle \notin \mathsf{Args}(\mathcal{L})$.

In Figure 2, we give the probability for each argument. Whilst this gives a perspective on the probability distribution, we have already acknowledged that we need to be selective in the arguments we present to better represent the probabilistic information.

2.4 Epistemic extensions

For a probability distribution P, and $A \in \mathsf{Args}(\mathcal{L})$, $P(A)$ is the degree of belief that A is acceptable. When $P(A) > 0.5$, then the argument is believed to be acceptable, whereas when $P(A) \leq 0.5$, then the argument is not believed to be acceptable.

Definition 7. *The* **epistemic extension** *for probability distribution* P *and graph* G *is* $\mathsf{Extension}(P, G) = \{A \in \mathsf{Args}(G) \mid P(A) > 0.5\}$.

Example 3. *For the exhaustive graph* G *in Figure 2, we get* $\mathsf{Extension}(P, G) = \{\langle\{\}, a\rangle, \langle\{\}, b\rangle, \langle\{b\}, a\rangle, \langle\{a\}, b\rangle\}$.

The epistemic approach provides a finer grained assessment of an argument graph than given by Dung's definition of extensions. By adopting constraints on the distribution, the epistemic approach subsumes Dung's approach [30, 17]. However, there is also a need for a non-standard view where we adopt alternative constraints on the distribution. For instance, we may wish to represent disbelief in arguments even when they are unattacked [22]. Nonetheless, for the non-standard view we may want the probabilities to respect the structure of the graph in some sense [12, 17, 22]. For example, when argument A attacks argument B, the **rational constraint** ensures that if A is believed (i.e. $P(A) > 0.5$), then B is not believed (i.e. $P(B) \leq 0.5$), and **coherence constraint** ensures that the belief in A and B sum to less than or equal to 1 (i.e. $P(A) + P(B) \leq 1$).

Example 4. *Examples of belief in arguments in Figure 3.*

	A_1	A_2	A_3	A_4	A_5	*Rational*	*Coherence*
P_1	*0.6*	*0.9*	*0.4*	*0.6*	*0.7*	*No*	*No*
P_2	*0.3*	*0.9*	*0.3*	*0.1*	*0.8*	*Yes*	*No*
P_3	*0.9*	*0.1*	*0.2*	*0.8*	*0.2*	*Yes*	*Yes*

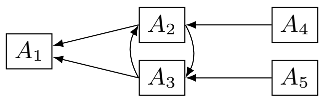

Figure 3: Example of argument graph.

So with constraints, we can manage how the structure of the graph is reflected in the probability distribution (and vice versa), and thereby it is a way to combine the belief in the composition of an argument, and the belief in acceptability of its counterarguments.

2.5 Plausibility of an argument

The plausibility measure captures the probability of the argument if the antecedent is subsequently found to be true. This is calculated by increasing the mass assigned to each model that satisfies the support of the argument in proportion to its current mass.

Definition 8. *The* **plausibility distribution***, denoted* $\hat{P}$*, for a probability distribution* P *and an argument* A *is calculated as follows, where* $\kappa = \sum_{X' \models \mathsf{Support}(A)} P(X')$*. If* $\kappa = 0$*, then* $\hat{P} = P$*, otherwise* $\hat{P}$ *is defined for each* $X \subseteq \mathcal{S}$ *as follow.*

$$\hat{P}(X) = \begin{cases} P(X)/\kappa & \textit{if } X \models \mathsf{Support}(A) \\ 0 & \textit{if } X \not\models \mathsf{Support}(A) \end{cases}$$

So the **plausibility of an argument** A *is*

$$\hat{P}(A) = \sum_{X \models \mathsf{Support}(A) \wedge \mathsf{Claim}(A)} \hat{P}(X)$$

Example 5. *Continuing Example 2, for* $A_1 = \langle\{b\}, f\rangle$, $P(A_1) = 0.95$ *and* $\hat{P}(A_1) = 0.95$ *and for* $A_2 = \langle\{p, b\}, \neg f\rangle$, $P(A_2) = 0.01$ *and* $\hat{P}(A_2) = 1$.

At the level of arguments, the plausibility function is probabilistic conditioning, and so similar to Pollock's argument strength [25]. Unlike Pollock, we do use the measure in an inference mechanism.

Proposition 5. *For all* A, $\hat{P}(A) = P(A)/P(\mathsf{Support}(A))$.

Therefore, for any probability distribution P and argument A, $P(A) \leq \hat{P}(A)$. The following results consider what happens when the probability and plausibility coincide, the effect of zero probability, and the nature of the uniform distribution.

Proposition 6. *For all* P, $P(A) = \hat{P}(A)$ *iff either for all* $X \subseteq \mathcal{S}$, *if* $P(X) \neq 0$, *then* $X \models \mathsf{Support}(A)$, *or for all* $X \subseteq \mathcal{S}$, *if* $P(X) \neq 0$, *then* $X \not\models \mathsf{Support}(A)$.

Proposition 7. *For all* P, *if* $P(A) = 0$, *then* $\hat{P}(A) = 0$.

Proposition 8. *For* n *atoms, let* G^* *be the exhaustive graph. If* P *is a uniform distribution, then for each* $A \in \mathsf{Nodes}(G^*)$, $1/2^n \leq P(A) \leq 1/2$ *and* $\hat{P}(A) = 1/2$.

In the next section, we will use the plausibility measure to help select arguments to present in an argument graph.

3 SELECTIVITY IN ARGUMENTATION

As we have seen from Figure 2, the exhaustive graph has an overwhelming number of arguments and attacks. To address this, we use the probability distribution to help select arguments and attacks to include in the argument graph, thereby providing a better representation by focusing on the key arguments and attacks.

3.1 Desiderata for selective graphs

Since, we assume the primacy of the probability distribution, the aim of the argument graph is to explain the distribution with particular emphasis on a query (i.e. a literal that is a claim of interest). We also assume explanations should adhere to some principles concerning what is shown and what is not selected. The following are desiderata where G is a selective argument graph, P is a probability distribution, and $\hat{P}$ is the corresponding plausibility function.

Faithful If G^* is the exhaustive graph, then $G \sqsubseteq G^*$.
Relevant If $\psi \in \mathcal{L}$ is a query, then there is an $A \in \mathsf{Nodes}(G)$ such that $\mathsf{Claim}(A) \in \{\psi, \neg\psi\}$.
Plausible For $(B, A) \in \mathsf{Arcs}(G)$, $\hat{P}(B) \geq 0.5$.
Informative For $(B, A) \in \mathsf{Arcs}(G)$, if B is not a source in G (i.e. there is an undercut for B), then $\mathsf{Atoms}(B) \neq \mathsf{Atoms}(A)$.
Rational For $(B, A) \in \mathsf{Arcs}(G)$, if $P(B) > 0.5$, then $P(A) \leq 0.5$.
Semi-optimistic If for every attacked argument A, $P(A) \geq 1 - \sum_{B \text{ s.t. } (B,A) \in \mathsf{Arcs}(G)} P(B)$.

We explain these properties as follows: (Faithful) selective graphs should only contain nodes and arcs from the exhaustive graph; (Relevant) the query should be an atom in the claim of an argument in the selective graph; (Plausible) for each attack in the selective graph, the attacker is plausible; (Informative) for each attack in the selective graph, the attacker and attackee do not just refer to the same atoms (unless the attacker is unattacked) – the aim is to avoid a chain of argument, undercut, undercut to undercut, where the same atoms are used, and so avoids unhelpful repetition of information; (Rational) for each attack, if the attacker is believed, then the attackee is not believed; (Semi-optimistic) for each attackee, there is a lower bound on belief in it (e.g. A is only attacked by B, and $P(B) = 0.2$, then $P(A) \geq 0.8$).

3.2 Constructing selective graphs

We now consider how to construct a selective graph that meets the desiderata. There are various ways we may wish to do this. We will consider two options in detail. We start with the following subsidiary definitions.

Definition 9. *Let* ψ *be a query where* $\psi \in \mathcal{S}$ *and let* G^* *be the exhaustive graph. A* **query argument** *for* ψ *is an argument* $A \in \mathsf{Nodes}(G^*)$ *such that* $\mathsf{Claim}(A) \in \{\psi, \neg\psi\}$. *A* **starting argument** *for* ψ *is a query argument* $A \in \mathsf{Nodes}(G^*)$ *for* ψ *such that* $P(A) \geq 0.5$ *and for all query arguments* $A' \in \mathsf{Nodes}(G^*)$ *for* ψ *such that* $P(A') \geq 0.5$, $\mathsf{Support}(A) \subseteq \mathsf{Support}(A')$.

In general, for a query ψ, probability distribution P, and exhaustive graph G^*, there is not a unique starting argument.

Example 6. *Consider* $\mathcal{S} = \{a, b\}$. *For query* a, *the query arguments are* $A_1 = \langle\{\}, a\rangle$, $A_2 = \langle\{\}, \neg a\rangle$ $A_3 = \langle\{b\}, a\rangle$, $A_4 = \langle\{\neg b\}, a\rangle$ $A_5 = \langle\{b\}, \neg a\rangle$, *and* $A_6 = \langle\{\neg b\}, \neg a\rangle$. *For the following probability distribution,* A_1, A_2, A_3, *and* A_6 *are starting arguments.*

$\langle a, b\rangle$	*11*	*10*	*01*	*00*
P	*0.5*	*0.0*	*0.0*	*0.5*

In the definition of a selective graph, we can use the notion of a **fixed set** $\Psi \subseteq \mathcal{L}$ to specify that all arguments that appear in the graph have a support that contains the fixed set. For example, returning to Example 2, we could have the fixed set $\{b\}$, and so all arguments A would be required to have $b \in \mathsf{Support}(A)$. Both definitions for selective graphs in this paper use this requirement but we are not proposing that any definition for a selective graph has to use it.

Our first type of selective graph is the procon window that presents the starting arguments with supports that contain the fixed set. This gives a bipartite graph.

Definition 10. *A* **procon window** *for a query* ϕ *and fixed set* Ψ *is a graph* $G \sqsubseteq G^*$ *s.t.* $\mathsf{Nodes}(G) = \{B \in \mathsf{Nodes}(G^*) \mid B$ *is a query argument for* ϕ *s.t.* $\hat{P}(B) \geq 0.5$ *and* $\Phi \subseteq \mathsf{Support}(B)\}$ *and* $\mathsf{Arcs}(G) = \{(B, A) \in \mathsf{Arcs}(G^*) \mid B, A \in \mathsf{Nodes}(G)\}$.

Note, in the examples of selective graphs in the rest of the paper, the probability (respectively plausibility) of the argument is given on the left (respectively right) in square brackets after the argument.

Example 7. *Continuing Example 6, with fixed set* $\Phi = \{\}$ *and query* a, *the following is the procon window.*

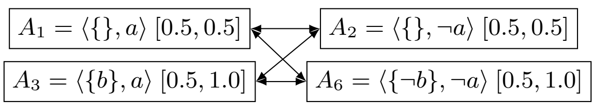

The following simple example indicates how procon windows can be used to capture useful arguments for and against a diagnosis. In this case, it is whether a particular symptom results in a particular disease. It would be straightforward to generalize to consider multiple symptoms and signs.

Example 8. *Consider atoms d denoting "*disease*" and s denoting "*symptom*", and the following probability distribution.*

$\langle d, s\rangle$	*11*	*10*	*01*	*00*
P	*0.09*	*0.01*	*0.0*	*0.9*

From P, the fixed set $\Phi = \emptyset$, and the query d, we get the following procon window showing strong belief in arguments against the disease, together with a highly plausible argument for the disease.

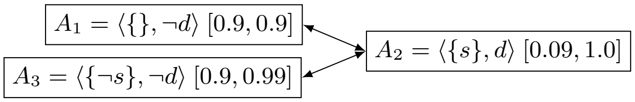

Our second type of selective graph is the span window which is a subgraph with a limit on the length of paths from each node to the starting node. For this, we require the following subsidiary definitions: For a graph G, and arguments $A, B \in \mathsf{Nodes}(G)$, the **pathlength** function, denoted $\mathsf{PathLength}(A, B)$, gives the length the shortest path from A to B (i.e. the lowest number of arcs to traverse from A to B). For example, in Figure 3, $\mathsf{PathLength}(A_4, A_1) = 2$ whereas $\mathsf{PathLength}(A_1, A_4) = \infty$. The **span** σ of a graph G with a starting argument A is the highest value for pathlength of the arguments $B \in \mathsf{Nodes}(G)$. So for example, for $\sigma = 0$, the graph has just the starting argument, for $\sigma = 1$, the graph has the starting argument and possibly direct attackers, for $\sigma = 2$, the graph has the starting argument, and possibly direct attackers, and possibly direct attackers to the direct attackers, and so on.

Definition 11. *A* **span window** *for a starting argument A, a fixed set $\Psi \subseteq \mathcal{L}$, and a span $\sigma \in \mathbb{N}$, is an acyclic graph $G \sqsubseteq G^*$ s.t.:*

1. *for all $B \in \mathsf{Nodes}(G)$, (a) $\mathsf{PathLength}(B, A) \leq \sigma$; (b) $\Psi \subseteq \mathsf{Support}(B)$; and (c) if $B \neq A$, then $\hat{P} > 0.5$.*
2. *for all $(B, C) \in \mathsf{Arcs}(G)$, (a) if B is not a presumption, then $\mathsf{Atoms}(B) \not\subseteq \mathsf{Atoms}(C)$; (b) if B is a presumption, then B is a source; and (c) $\mathsf{PathLength}(B, A) \geq \mathsf{PathLength}(C, A)$*

We explain the conditions as follows: (1a) The span of the graph is bounded by σ; (1b) The fixed set appears in the support of each argument in the graph; (1c) Apart from possibly the starting argument, each argument has a plausibility greater than 0.5; (2a) Each attacker has some atoms different to its attackee apart from when it is a presumption (so it is a specificity principle that filters out attackers that have more general support) – e.g. in Example 9 attack by $\langle\{b\}, f\rangle$ on $\langle\{p, b\}, \neg f\rangle$ is not possible, whereas from $\langle\{p, b\}, \neg f\rangle to \langle\{b\}, f\rangle$ is possible; (2b) Apart from possibly the starting argument, only a source (i.e. an unattacked argument) has the possibility of being presumptive; and (2c) Each path goes to the starting argument.

Example 9. *Returning to Example 2, let the query be f, we get the query arguments $\langle\{b\}, f\rangle$ and $\langle\{\neg p, b\}, f\rangle$ that are believed. Of these, $\langle\{\neg p, b\}, f\rangle$ is the only starting argument.*

$\langle b, p, f\rangle$	*110*	*101*	*100*
P	*0.01*	*0.95*	*0.04*

From P, and the fixed set $\emptyset$, we obtain the following span window for $\sigma \geq 3$. Whereas if $\sigma = 2$, then we would not include $\langle\{\}, \neg p\rangle$ or $\langle\{b\}, \neg p\rangle$. Note, $\langle\{b\}, \neg p\rangle$ does not attack $\langle\{p, b\}, \neg f\rangle$ because of condition 2a from Definition 11.

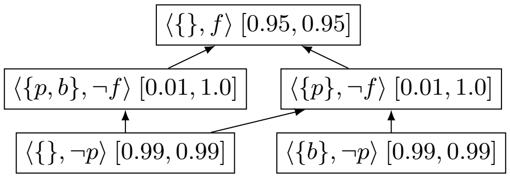

Furthermore, with P and the fixed set $\{b\}$, we obtain the following span window for $\sigma \geq 2$.

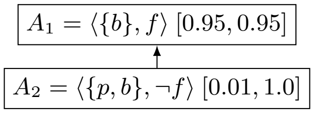

Example 10. *Let c denote "*patient has a cold*", f denote "*patient has flu*", and e denote "*there is a flu epidemic currently*". Consider the following probability distributions.*

$\langle c, f, e\rangle$	*101*	*100*	*011*	*010*	*001*	*000*
P	*0.05*	*0.7*	*0.1*	*0.02*	*0.03*	*0.1*

From P, we obtain the following span window for $\sigma \geq 3$. So by presumption, the patient does not have a cold. Whilst, there are plausible counterarguments, there are strong reasons to reject those.

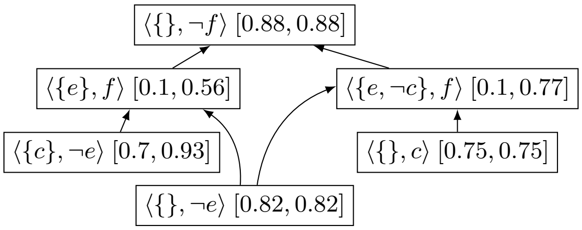

Example 11. *We extend Example 10 with the atom s to denote "*patient displays symptoms of flu*", and the following probability distribution.*

$\langle s, c, f, e\rangle$	*1101*	*1011*	*1001*
P	*0.1*	*0.8*	*0.1*

With P and $\Phi = \{s, e\}$, we get the following span window for $\sigma \geq 2$. So there is a strong argument for the patient having flu.

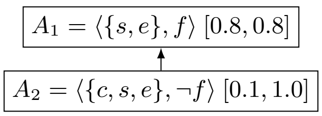

We now consider the graphs produced by the procon and span window definitions with respect to the desiderata properties.

Proposition 9. *If G is the selective graph for G^* obtained as a procon (respectively span) window, then G satisfies the properties of faithful, relevant, plausible, informative, rational, and semi-optimistic, given at the start of this subsection.*

The epistemic extension of a selective graph G is sound with respect to the epistemic extension of the exhaustive graph G^*.

Proposition 10. *For a probability distribution P, exhaustive graph G^*, and selective graph $G \sqsubseteq G^*$,*

$$\mathsf{Extension}(P, G) = \mathsf{Extension}(P, G^*) \cap \mathsf{Nodes}(G)$$

Clearly there are various options for windows from an exhaustive graph. We have provided two definitions that meet a set of desirable properties. However, for different applications, we may choose an alternative set of desirable properties. For instance, for more literals, we generate more arguments, and we may want to restrict the length of branches by limiting the support of arguments to have at least one atom that has not occurred on the branch to the starting argument or we may want to restrict the branching factor by limiting the overlap between siblings. These longer range requirements can then be reflected in new definitions for windows. We may also consider alternatives to the plausibility measure for deciding which arguments to include in the window.

4 COMMITMENT IN ARGUMENTATION

We now consider how we can analyse an argument graph by updating a probability distribution by making commitments to specific atomic propositions which we do by using probability statements. A probability statement is of the form (α, v) where $\alpha \in \mathcal{S}$ is an atomic proposition and $v \in [0, 1]$.

Definition 12. *Let Γ be a set of probability statements. A probability distribution P is* **committed** *to Γ iff for each $(\alpha, v) \in \Gamma$, $P(\alpha) = v$. A set of probability statements Γ is* **consistent** *iff there is a probability distribution P such that P is committed to Γ.*

Commitment is not necessarily to 0 or 1 because there may still be uncertainty about the atom.

Definition 13. *Let P be a probability distribution, and let Γ be a consistent set of probability statements (i.e. if $p(\alpha, v_1) \in \Gamma$, and $(\alpha, v_2) \in \Gamma$, then $v_1 = v_2$). A* **commitment function** *is function of the form $\mathsf{Com}(P, \Gamma) = P'$.*

Since argumentation is about dealing with incomplete and inconsistent information, commitment can be used as part of a process of eliminating uncertainty. At the start perhaps there is uncertainty about all the atoms in $\mathcal{S}$, and then as the problem is investigated (perhaps by asking questions, doing research, seeking advice, etc), commitments are made to atoms

Commitment can also be used hypothetically. In this way, commitments can be made to investigate scenarios. So different choices for Γ would denote different scenarios. From this, the robustness of specific arguments could be identified. For instance, if a particular argument, or a particular claim, occurs for a wide variety of choices of Γ, then that argument or claim would appear to be more robust.

The following are some optional requirements for commitment: (Tautology) An empty commitment should not change the probability distribution; (Idempotence) Repeating a commitment should not change the result; and (Invariance) The order in which the commitments are done should not affect the result.

Tautology $\mathsf{Com}(P, \emptyset) = P$.
Idempotence If $\mathsf{Com}(P, \Gamma) = P'$ then $\mathsf{Com}(P', \Gamma) = P'$.
Invariance $\mathsf{Com}(\mathsf{Com}(P, \Gamma_1), \Gamma_2) = \mathsf{Com}(\mathsf{Com}(P, \Gamma_2), \Gamma_1)$.

Next, we define a particular commitment function which updates the satisfying models by the ratio of the desired and current value for the atom, and similarly for the non-satisfying models. For this, we assume that for any commitment (α, v), $0 < P(\alpha) < 1$ (i.e. it is not the case that α is completely believed or disbelieved). It is straightforward to generalize the definition in order to drop this restriction.

Definition 14. *A* **simple commitment** *by probability distribution P to a statement (α, v), where $P(\alpha) \in (0, 1)$, is a distribution P', denoted $\mathsf{Simple}(P, (\alpha, v))$, as follows, where $X \subseteq \mathcal{S}$ is a model:*

$$P'(X) = \begin{cases} P(X) \times \frac{v}{P(\alpha)} & \text{if } X \models \alpha \text{ and } v > 0 \\ P(X) \times \frac{1-v}{1-P(\alpha)} & \text{if } X \not\models \alpha \text{ and } v < 1 \\ 0 & \text{otherwise} \end{cases}$$

For $\Gamma = \{(\alpha_1, v_1), \ldots, (\alpha_n, v_n)\}$, the simple commitment function, denoted $\mathsf{Com}_s(P, \Gamma) = P_n$, is defined recursively as follows for $i \in \{1, \ldots, n\}$, where P_0 is P, and $\Gamma_{i+1} = \Gamma_i \setminus \{(\alpha_i, v_i)\}$.

$$\mathsf{Com}_s(P_{i-1}, \{(\alpha_i, v_i)\} \cup \Gamma_{i+1}) = \mathsf{Com}_s(\mathsf{Simple}(P_{i-1}, (\alpha_i, v_i)), \Gamma_{i+1})$$
$$\mathsf{Com}_s(P_n, \{\}) = P_n$$

So we obtain P_1 from P_0 and (α_1, v_1), P_2 from P_1 and (α_2, v_2), and so on until P_n from P_{n-1} and (α_n, v_n).

Example 12. *Below we give a probability distribution P, and the result distribution for each of three alternatives for committing to b.*

Update	$\langle a, b \rangle$	*11*	*10*	*01*	*00*
	P	*0.05*	*0.75*	*0.1*	*0.1*
P_1	$(b, 1)$	*0.33*	*0*	*0.67*	*0*
P_2	$(b, 0.4)$	*0.13*	*0.53*	*0.27*	*0.07*
P_3	$(b, 0.1)$	*0.03*	*0.79*	*0.07*	*0.11*

Commitment causes mass to be transferred to a smaller subset of models.

Example 13. *Continuing with the atoms in Example 2, we start with P_1 being all instances of vertebrate in a particular zoo. Using commitment, P_2 is obtained from P_1 with the commitment $(b, 1)$ (i.e. it is a bird), and then P_3 is obtained from P_2 with the commitment $(p, 0)$ (i.e. it is not a penguin),*

$\langle b, p, f \rangle$	*110*	*101*	*100*	*000*
P_1	*0.01*	*0.085*	*0.005*	*0.9*
P_2	*0.1*	*0.85*	*0.05*	*0*
P_3	*0*	*0.94*	*0.06*	*0*

From P_1, we obtain the following span window for $\sigma \geq 3$.

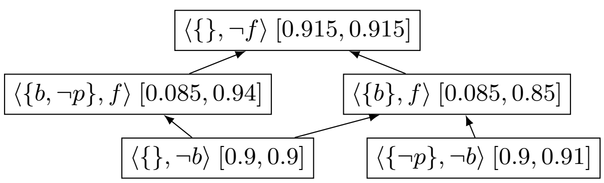

From P_2, we obtain the following with $\sigma \geq 2$ and $\Phi = \{b\}$.

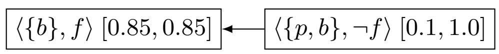

From P_3, we obtain the following with $\sigma \geq 1$ and $\Phi = \{b, \neg p\}$.

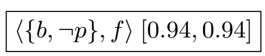

The simple commitment function produces a probability distribution.

Proposition 11. *For a probability distribution P, and statements Γ, if* $\mathsf{Com}_s(P, \Gamma) = P'$*, then P' is a probability distribution.*

For any commitments, the simple commitment function satisfies tautology and idempotence. However, when they are restricted to 0 or 1, then it also satisfies invariance.

Proposition 12. *For $\Gamma = \{(\alpha_1, v_1), \ldots, (\alpha_n, v_n)\}$, if each $i \in \{1, \ldots, n\}$, it is the case that $v_i \in \{0, 1\}$, then* $\mathsf{Com}_s(P, \Gamma)$ *function satisfies tautology, idempotence, and invariance.*

If a probability distribution already agrees with a probability statement, then there is no change with simple commitment.

Proposition 13. *If $P(\alpha) = v$ and* $\mathsf{Com}_s(P, \{(\alpha, v)\}) = P'$*, then $P = P'$.*

In the following, we show that the simple commitment function is a generalization of the plausibility measure and therefore of probabilistic conditioning. In the plausibility measure, there is a hypothetical commitment to the support of an argument being completely believed, whereas commitment is to any literal, and the commitment can be for any value in the unit interval (as long we do not require invariance).

Proposition 14. *For an argument A, if* $\mathsf{Support}(A) = \{\beta_1, \ldots, \beta_n\}$*, and* $P(\mathsf{Support}(A)) > 0$*, then $\hat{P}(A) = P'(A)$, where $\Gamma = \{(\beta_1, 1), \ldots, (\beta_n, 1)\}$, and* $\mathsf{Com}_s(P, \Gamma) = P'$*.*

Finally, we formalize an example of defeasible reasoning taken from Pollock [24] concerning whether something is red, thereby showing how the updates can be captured as commitments.

Example 14. *Let i denote* "illumination is red"*, o denote* "looks red"*, and r denote* "it is red"*. Consider query r, and fixed set $\{o\}$.*

$\langle i, o, r\rangle$	*111*	*110*	*011*
P	0.025	0.075	0.9
P'	0.25	0.75	0

From P, we obtain the following span window for $\sigma \geq 2$.

$$\boxed{A_1 = \langle \{o\}, r\rangle\ [0.93, 0.93]}$$
$$\uparrow$$
$$\boxed{A_2 = \langle \{i, o\}, \neg r\rangle\ [0.07, 0.75]}$$

So argument A_1, which is believed, says that with the observation of the item looking red, the item is red, whereas the counterargument, which is disbelieved, says that with the observation of the item looking red, together with the illumination being red, the item is not red. So the counterargument is not believed, but the plausibility is high, indicating a reason to doubt the argument A_1. Furthermore, if the probability distribution is revised to give belief in $P'(i) = 1$ and $P'(o) = 1$, and the fixed set is $\Phi = \{i, o\}$, then the selective graph is the following with just one node. So now the argument against the item being red is believed.

$$\boxed{A_2' = \langle \{i, o\}, \neg r\rangle\ [0.75, 0.75]}$$

We propose the simple commitment function as an illustration of how we can define commitment, and show how it can be useful to capture updates that are not necessarily categorical. Alternative approaches to defining commitment functions include Jeffrey's updating rule [18], refinement [13], and distance-based methods that minimize the change to a probability distribution [16].

5 COMPARISON WITH LITERATURE

The two main approaches to probabilistic (abstract) argumentation are the constellations and the epistemic approaches [12]. In the epistemic approach, the topology of the argument graph is fixed, but there is uncertainty about whether an argument is believed [30, 12, 1, 13, 7, 14, 15, 17]. The epistemic approach has been extended to also allow a probability distribution over subsets of attacks, and thereby represent belief in each attack [23].

At the structured level, Haenni [10] considered a restricted form of probabilistic argumentation in which pros and cons are generated from a classical logic knowledgebase, and then a probability distribution over models of the language are used to assign a belief in each argument. Subsequently, this was generalized by Hunter to arbitrary argument graphs [12] in which various kinds of counterargument can be accommodated. More recently, Prakken considered a similar approach for ASPIC+ [26]. In other logic-based proposals, Verheij has combined probabilities with non-monotonic inference [32] and separately, he has combined qualitative reasoning in terms of reasons and defeaters (adapting Pollock's definitions [25]), with quantitative reasoning using argument strength (modeled as the conditional probability of the conclusions given the premises) [33]. In these five proposals, it is assumed that somehow a knowledgebase of formulae is available, that the arguments are constructed from this knowledgebase, and that the probability values are then assigned to the arguments. In contrast, in this paper, we generate all possible arguments based in the language, then we select arguments for the graph based on the probability distribution and requirements for the graph.

Dung and Thang [5] provided the first proposal for a probability distribution over sets of arguments which is used to obtain a probability distribution over induced subgraphs. This is then used with a probabilistic version of assumption-based argumentation. Then, Li *et al* [19] proposed a probability assignment to arguments and attacks, and when assuming independence, these can be used to generate a probability distribution over subgraphs. Both these proposals can be viewed as special cases of the constellation approach (as opposed to the epistemic approach used in this paper). In another rule-based system for dialogical argumentation, the belief in the premises of an argument is used to calculate the belief in the argument, though the nature of this belief is not investigated [29].

Bayesian networks can be used to model argumentative reasoning with arguments and counterarguments [34]. 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 [9]. Argumentation can also be used to help construct Bayesian networks [2, 35]. Going the other way, arguments can be generated from a Bayesian network, and this can be used to explain the Bayesian network [31]. This involves constructing arguments involving a rule-based language in ASPIC+ for reflecting the network structure. Finally, argumentation can be used to combine multiple Bayesian networks [20]. However, none of the above works offer a way to generate the arguments from the probability distribution, or a way to select arguments to present in a graph, as we have provided in this paper.

6 DISCUSSION

In this paper, we have provided a framework for generating and analyzing argument graphs from a probability distribution. This offers a more accessible and better explained way of viewing and analysing the distribution (as illustrated in the examples, each graph is easier to assimilate than its corresponding probability distribution). Furthermore, it is configurable. We can select the choice of arguments/attacks as a type of window, and we can consider the requirements of the window as a set of constraints. We can consider different requirements as desirable properties, and we can draw on measures such as belief and plausibility to define these requirements and windows. This proposal is interesting for applications as we can obtain a probability distribution from data, ML models, or asking people for belief in scenarios. Argumentation can then make sense of the distribution, provide explanations, or be used to make decisions.

Being selective when constructing argument graphs appears to reflect how humans are not exhaustive when arguing. This selectivity may be based on various criteria, but belief in arguments is likely to be an important aspect. In future work, we will investigate this in empirical evaluation with participants, and so follow a recent trend in evaluating argumentation formalisms [27, 4, 22]. This may include studying whether people prefer some kinds of window over others, whether given the probabilistic information, people would tend to produce the same arguments and counterarguments as our approach, and whether when people are given an argument graph produced by our methods, they understand the probabilistic information better (perhaps by providing better answers to questions). Also in future work, we will consider situations without a unique probability distribution. For example, we might only know the belief in some formulae which might be the case if we are dealing with enthymemes. For this, we may use the principle of maximum entropy to select a unique probability distribution (see for example [21, 17]).

REFERENCES

[1] P. Baroni, M. Giacomin, and P. Vicig, 'On rationality conditions for epistemic probabilities in abstract argumentation', in *Proceedings of the Fifth International Conference on Computational Models of Argumentation (COMMA'14)*, (2014).

[2] F. Bex and S. Renooij, 'From arguments to constraints on a Bayesian network', in *Proceedings of the Sixth International Conference on Computational Models of Argumentation (COMMA'16)*, (2016).

[3] M. Caminada and D. Gabbay, 'A logical account of formal argumentation', *Studia Logica*, **93**, 109–145, (2009).

[4] F. Cerutti, N. Tintarev, and N. Oren, 'Formal arguments, preferences, and natural language interfaces to humans: an empirical evaluation', in *Proceedings of the European Conference on Artificial Intelligence (ECAI'14)*, volume 263 of *FAIA*, pp. 1033–1034. IOS Press, (2014).

[5] P. Dung and P. Thang, 'Towards (probabilistic) argumentation for jury-based dispute resolution', in *Proceedings of the Third International Conference on Computational Models of Argumentation (COMMA'10)*, pp. 171–182. IOS Press, (2010).

[6] 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), 321–358, (1995).

[7] D. Gabbay and O. Rodrigues, 'Probabilistic argumentation: an equational approach', *Logica Universalis*, **9**(3), 345–382, (2015).

[8] N. Gorogiannis and A. Hunter, 'Instantiating abstract argumentation with classical logic arguments: Postulates and properties', *Artificial Intelligence*, **175**(9-10), 1479–1497, (2011).

[9] M. Grabmair, T. Gordon, and D. Walton, 'Probabilistic semantics for the carneades argument model using Bayesian networks', in *Proceedings of the Third International Conference on Computational Models of Argumentation (COMMA'10)*, pp. 255–266, (2010).

[10] R. Haenni, 'Modelling uncertainty with propositional assumptions-based systems', in *Applications of Uncertainty Formalisms*, LNCS, 446–470, Springer, (1998).

[11] A. Hunter, 'Some foundations for probabilistic abstract argumentation', in *Proceedings of the Fourth International Conference on Computational Models of Argumentation (COMMA'12)*, pp. 117–128. IOS Press, (2012).

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

[13] A. Hunter, 'Modelling the persuadee in asymmetric argumentation dialogues for persuasion', in *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI'15)*, (2015).

[14] A. Hunter, 'Persuasion dialogues via restricted interfaces using probabilistic argumentation', in *Proceedings of the 10th International Conference on Scalable Uncertainty Management (SUM'16)*, (2016).

[15] A. Hunter, 'Two dimensional uncertainty in persuadee modelling in argumentation', in *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI'16)*, (2016).

[16] A. Hunter and N. Potyka, 'Updating probabilistic epistemic states in persuasion dialogues,', in *Proceedings of European Conference on Symbolic and Quantitative Approaches to Uncertainty (ECSQARU'17)*, volume 10369 of *LNCS*, pp. 46–56, (2017).

[17] A Hunter and M Thimm, 'Probabilistic reasoning with abstract argumentation frameworks', *Journal of Artificial Intelligence Research*, **59**, 565–611, (2017).

[18] R. Jeffrey, *The Logic of Decision*, Chicago University Press, 1965.

[19] H. Li, N. Oren, and T. J. Norman, 'Probabilistic argumentation frameworks', in *Proceedings of the First International Workshop on the Theory and Applications of Formal Argumentation (TAFA'11)*, (2011).

[20] S. Nielsen and S. Parsons, 'An application of formal argumentation: fusing Bayesian networks in multi-agent systems', *Artificial Intelligence*, **171**(10-15), 754–775, (2007).

[21] N. Nilsson, 'Probabilistic logic', *Artificial Intelligence*, **28**, 71–87, (1995).

[22] S Polberg and A Hunter, 'Empirical evaluation of abstract argumentation: Supporting the need for bipolar and probabilistic approaches', *International Journal of Approximate Reasoning*, **93**, 487–543, (2018).

[23] S. Polberg, A. Hunter, and M. Thimm, 'Belief in attacks in epistemic probabilistic argumentation', in *Proceedings of the International Conference on Scalable Uncertainty Management (SUM'17)*, volume 10564 of *LNCS*, pp. 223–236. Springer, (2017).

[24] J. Pollock, 'Defeasible reasoning', *Cognitive Science*, **11**(4), 481–518, (1987).

[25] J. Pollock, *Cognitve Carpentry*, MIT Press, 1995.

[26] H. Prakken, 'Probabilistic strength of arguments with structure', in *Proceedings of International Conference on Principles of Knowledge Representation and Reasoning (KR'18)*, (2018).

[27] I. Rahwan, M Madakkatel, J. Bonnefon, R Awan, and S. Abdallah, 'Behavioural experiments for assessing the abstract argumentation semantics of reinstatement', *Cognitive Science*, **34**(8), 1483–1502, (2010).

[28] R. Riveret and G. Governatori, 'On learning attacks in probabilistic abstract argumentation', in *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'16)*, (2016).

[29] R. Riveret, A. Rotolo, G. Sartor, H. Prakken, and B. Roth, 'Success chances in argument games: a probabilistic approach to legal disputes', in *Proceedings of the 20th Annual Conference on Legal Knowledge and Information Systems (JURIX'07)*, pp. 99–108. IOS Press, (2007).

[30] M. Thimm, 'A probabilistic semantics for abstract argumentation', in *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI'12)*, (August 2012).

[31] S. T. Timmer, J.-J. Ch. Meyer, H. Prakken, S. Renooij, and B. Verheij, 'Explaining Bayesian networks using argumentation', in *Proceedings of the 13th European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'15)*, (2015).

[32] B. Verheij, 'Jumping to conclusions: A logico-probabilistic foundation for defeasible rule-based arguments', in *Proceedings of the European Conference on Logics in Artificial Intelligence (JELIA'12)*, volume 7519 of *LNCS*, pp. 411–423. Springer, (2012).

[33] B. Verheij, '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)*, (2014).

[34] G. Vreeswijk, 'Argumentation in Bayesian belief networks', in *Proceedings of the First International Workshop on Argumentation in Multi-Agent Systems (ArgMAS'04)*, pp. 111–129, (2004).

[35] R. Wieten, F. Bex, H. Prakken, and S. Renooij, 'Supporting discussions about forensic bayesian networks using argumentation', in *Proceedings of the International Conference on Artificial Intelligence and Law (ICAIL'19)*, pp. 143–152, (2019).