Chapter 7

Probabilistic Argumentation: A Survey

Anthony Hunter
Department of Computer Science, University College London, United Kingdom
anthony.hunter@ucl.ac.uk

Sylwia Polberg
School of Computer Science and Informatics, Cardiff University, United Kingdom
polbergs@cardiff.ac.uk

Nico Potyka
Institute for Parallel and Distributed Systems, University of Stuttgart, Germany
nico.potyka@ipvs.uni-stuttgart.de

Tjitze Rienstra
Institute for Web Science and Technologies, University of Koblenz-Landau, Germany
rienstra@uni-koblenz.de

Matthias Thimm
Institute for Web Science and Technologies, University of Koblenz-Landau, Germany
thimm@uni-koblenz.de

Abstract

Argumentation is inherently pervaded by uncertainty, which can arise as a result of the context in which argumentation is used, the kinds of agents that are involved in a given situation, the types of arguments that are used, and more. One of the prominent approaches for handling uncertainty in argumentation is probabilistic argumentation, which offers means of quantifying the level of uncertainty we are dealing with. This chapter offers an overview of the state-of-the-art research in this area.

1 Introduction

Argumentation is inherently pervaded by uncertainty. One of the core concepts of defeasible reasoning and therefore argumentation is the fallibility of human perception, which forces us to be able to reason even with incomplete information and to be prepared to retract our conclusions in the face of new data. This is further compounded by applying argumentation in real-life situations, where uncertainty can arise as a result of the context in which argumentation is used, the kinds of agents that are involved in a given situation, the types of arguments that are used, and more. Thus, just like there are multiple sources of uncertainty in argumentation, there are multiple proposals in the literature towards modelling this kind of reasoning. One of the prominent approaches is *probabilistic argumentation*, which often offers means of quantifying the level of uncertainty we are dealing with. While using probability theory as a means to model uncertain aspects of argumentation has been questioned even by Pollock [1995], recent developments in the field showed its adequacy in practical matters and from the perspective of artificial discussion, see also [Verheij, 2014] for a discussion on this. This chapter offers an overview of the state-of-the-art research in this area, and we start by considering various examples that can benefit from incorporating probabilities.

Example 1.1 (Taken from [Hunter, 2014])**.** *Suppose that there are two witnesses to a criminal escaping in a car. The first witness says that the getaway car is red, and the other witness says that the getaway car is orange. If we take a strict interpretation of the colours, then we have two arguments a and b below, where each argument attacks the other.*

- a = *"The getaway car is red".*
- b = *"The getaway car is orange".*

For these arguments, it may be inappropriate to treat "red" and "orange" as contradictory. There is some ambiguity, and hence some imprecision, in the use of these terms. And so, it may be possible to regard these two terms as consistent together. So if we consider the argument graph, there is some uncertainty as to whether a attacks b and vice versa.

The above example highlights that uncertainty may arise when there is ambiguity, a form of imprecision in the language used in the arguments. Another important reason for uncertainty is that real-world arguments presented in natural language are normally enthymemes [Walton, 1989], i. e. arguments with a support that is insufficient for the claim to be entailed and/or a claim that is incomplete. This means that a given enthymeme could be completed into a full argument in more ways than one, and every interpretation may have a certain probability of being the intended one. We consider this in the next example.

Example 1.2 (Taken from [Hunter, 2014]). *Consider the following arguments:*

- a = *"The sun is shining now, we should organize a BBQ for this evening".*
- b = *"The weather report predicts rain this evening".*

Here, a is an argument that has incomplete premises for obtaining the claim (i. e. the premise "The sun is shining now" is insufficient for entailing the claim "We should organize a BBQ for this evening"). And b is an argument that has the premise "The weather report predicts rain this evening", but lacks a claim. Implicitly, the claim of b should negate the premises or the claim of argument a.

When a counterargument is an enthymeme, there may be uncertainty as to whether the argument being attacked is attacked because a premise is being contradicted or because the claim is being contradicted. In the above example, it could be that the implicit claim of b is negating the premise "The sun is shining now", or negating the claim "We should *organize a BBQ for this evening". Also, because a is an enthymeme with incomplete premises, b could have a claim that contradicts a missing premise of a. For instance, suppose we make the premises of argument a explicit as follows:*

- a_1 = *"The sun is shining now".*
- a_2 = *"If the sun is shining now, it will be warm and dry this evening".*
- a_3 = *"If it is warm and dry this evening, then we should organize a BBQ for this evening".*

So if the argument is explicit, it will have premises a_1, a_2 and a_3, and the claim "We should organize a BBQ for this evening". Then the claim of b can negate some combination of the explicit premises and/or the claim of a. Similarly, we can make the argument b explicit.

- b_1 = *"If the weather report predicts rain this evening, then it will not be warm and dry this evening".*
- b_2 = *"It will not be warm and dry this evening".*

Here b_1 is used to make the premises explicit, and b_2 is used to make the claim explicit. This claim then explicitly contradicts the premises a_1 and a_2. In other words, if we regard the premises of a as being represented by a_1, a_2 and a_3, and the premises of b being represented by b and b_1, and the claim of b being b_2, then we get the argument graph visible in Figure 1a.

However, there may be other interpretations of a and/or b, such that the interpretation of b does not attack the interpretation of a, then we get the argument graph in Figure 1b. In this way, there is doubt about whether a does indeed attack b, and what the precise structure of a and b is. All interpretations of the arguments and the resulting graphs are only given with some likelihood.

Yet another kind of uncertainty concerns the degree to which a given agent believes or disbelieves an argument or its premises or claims. In real-life, people often tend to trust or agree with certain things only up to

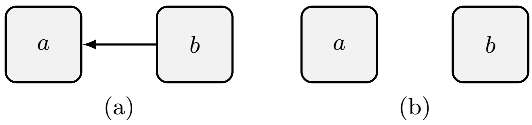

Figure 1: Possible argument graphs created from enthymemes.

a given degree, and being presented with counterarguments can weaken this belief rather than lead to outright rejection of a given argument. For instance, we might be inclined to have some doubt in the weather forecasts or in the stories of a gossiping neighbour. The way a given piece of information is presented to us also affects how we react to it, e. g., framing things a certain way or using particular language are classical strategies used in many areas, from marketing to healthcare. Finally, there is also the issue of imperfection. Human agents do not need to be perfect reasoners and can reach decisions in a biased or flawed manner; at the same time, we may accidentally deem them imperfect or not following certain rules of reasoning because they do not disclose all of their arguments and knowledge that would justify their position. Below we provide examples based on the empirical study from [Polberg and Hunter, 2018] that highlight some of the mentioned behaviours.

Example 1.3 (Taken from [Polberg and Hunter, 2018])**.** *Consider two listeners to the following discussion on flu shots between agents a and b.*

- a_1 = *Hospital staff members do not need to receive flu shots.*
- b_2 = *Hospital staff members are exposed to the flu virus a lot. Therefore, it would be good for them to receive flu shots in order to stay healthy.*
- a_3 = *The virus is only airborne and it is sufficient to wear a mask in order to protect yourself. Therefore, a vaccination is not necessary.*
- b_4 = *The flu virus is not just airborne, it can be transmitted through touch as well. Hence, a mask is insufficient to protect yourself against the virus.*

- a_5 = *The flu vaccine causes flu in order to gain immunity. Making people sick, who otherwise might have stayed healthy, is unreasonable.*

- b_6 = *The flu vaccine does not cause flu. It only has some side effects, such as headaches, that can be mistaken for flu symptoms.*

When asked, the first listener identifies certain conflicts between what the agents are uttering, as visible in Figure 2a. She admits to strongly agreeing with everything that agent b has said, however, she does not think that what agent a says is entirely wrong. She still, though to a small degree, agrees with a_1, *and* a_3 *prompts only slight disagreement (*a_5 *is completely disbelieved). She explains that while it is indeed a great idea to get vaccinated and it would be very beneficial to do so, she does not think that "needing" to do that is appropriate.*

In addition to the conflicts identified by the first listener, the second listener believes that a_5 *attacks* b_4 *and* b_6 *attacks* a_3 *(see Figure 2b). However, she holds a clear anti-vaccine stance and agrees (resp. disagrees) with everything agent a (resp. b) says. The listener states that agent b is expressing lies and inaccuracies and rejects the agent's views without providing any particular kind of evidence to the contrary, thus exhibiting behaviour that could be deemed as biased or not rational.*

The aforementioned examples highlight various kinds of uncertainty that can arise concerning the structure of the argument graph or the degree to which a given argument is accepted. While they are fairly straightforward, they are also quite common and reflect the ambiguity and uncertainty that is present in daily life and communication. They also represent challenges that any application of argumentation, particularly one involving human agents, will need to tackle.

The kinds of scenarios we have considered can be conveniently modelled using probabilistic argumentation, particularly with the constellations and the epistemic approaches [Hunter, 2013]. In the *constellations approach*, the uncertainty is in the topology of the graph, by which we understand that certain arguments or relations appear in the graph only with a given probability. This approach is useful when one agent is not sure what arguments and attacks another agent is aware of, or if ambiguity or imprecision of the arguments causes uncertainty in the structure

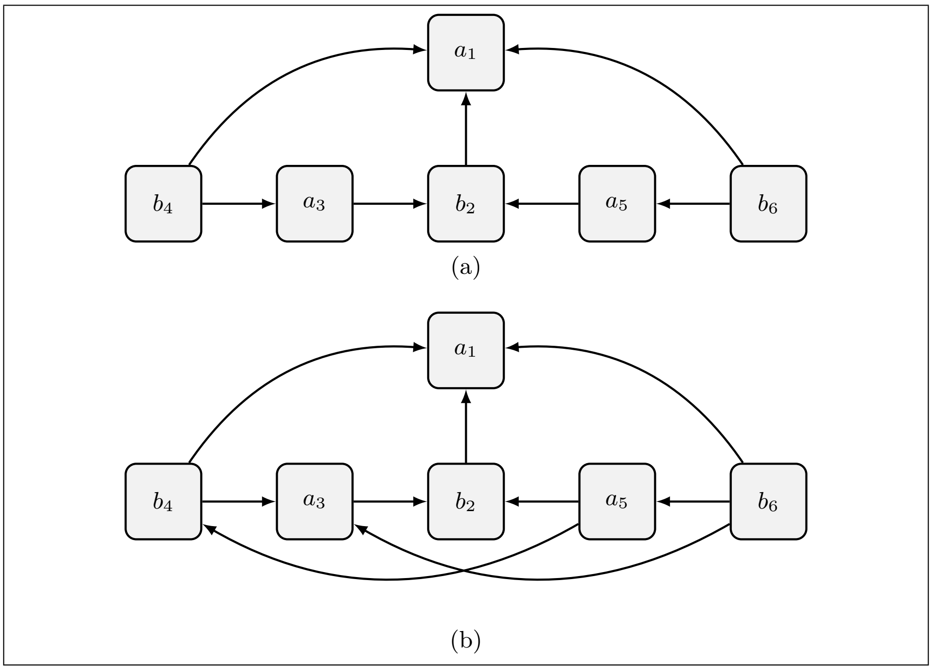

Figure 2: Example of attacks between arguments acquired by two listeners of a vaccine discussion.

of the graph. In the *epistemic approach*, the topology of the argument graph is fixed, but there is uncertainty as to the degree to which each argument is believed. This method can be harnessed when an agent is not certain of their own or another agent's opinion on the arguments, if a given situation calls for a fine-grained way to judge arguments, or if the perceived reasoning of a given agent escapes classical argumentation semantics. In this chapter we will discuss these and further kinds of probabilistic argumentation in more detail and explain how our examples can be modelled.

In Section 2 we recall certain basic notions that we will use throughout this chapter. Sections 3 and 4 investigate the epistemic and constellation approach to probabilistic argumentation, respectively, and Section 5 is dedicated to discussing other kinds of probabilistic approaches. We close with additional discussion and conclusions in Section 6.

2 Preliminaries

Within argumentation we frequently distinguish between the *structured* (or logic-based) and *abstract* approaches. The former define the relations between arguments in terms of the assumed internal structure of these arguments, and often offer means of constructing them from an underlying knowledge base or data source [Besnard and Hunter, 2008; Modgil and Prakken, 2014; Toni, 2014; Garcia and Simari, 2004]. In contrast, *abstract argumentation* [Dung, 1995; Brewka *et al.*, 2014] takes a very simple view on argumentation as it does not presuppose any internal structure of an argument. The arguments and relations between them are assumed to have been constructed, and the focus is put on determining what arguments can be deemed acceptable or not based on how they interact with each other. Various negative and positive kinds of relations have been studied in the literature [Brewka *et al.*, 2014], however, we will focus on the original argumentation framework by Dung [1995] that considers only a binary attack relation.

Definition 2.1. *An* abstract argumentation framework AF *is a tuple* $AF = (Ar, att)$ *where* Ar *is a set of arguments and* att *is a relation* $att \subseteq Ar \times Ar$.

Let $\mathbb{A}$ denote the set of all abstract argumentation frameworks. For two arguments $a, b \in Ar$ the relation $(a, b) \in att$ means that argument a attacks argument b. With $b^- = \{a \mid (a, b) \in att\}$ we denote the set of attackers of an argument b. For convenience, with $S^- = \bigcup_{a \in S} a^-$ we will denote the set of all attackers of a set of arguments $S \subseteq Ar$.

Abstract argumentation frameworks can be concisely represented by directed graphs, where arguments are represented as nodes and edges model the attack relation (see Figure 2)[1]. The status of a given argument is determined through the appropriate *argumentation semantics*, which often produce answers in the form of extensions [Dung, 1995] or labellings [Wu and Caminada, 2010]. In this work, we use the latter, though we note that for Dung's frameworks, they can be used interchangeably.

[1]Note that we only consider finite argumentation frameworks here, i. e., argumentation frameworks with a finite number of arguments.

Definition 2.2. *A* labelling $\mathcal{L}ab$ *for an abstract argumentation framework* $AF = (Ar, att)$ *is a function* $\mathcal{L}ab : Ar \to \{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$*.*

A labeling $\mathcal{L}ab$ assigns to each argument $a \in Ar$ either the value `in`, meaning that the argument is accepted, `out`, meaning that the argument is rejected, or `undec`, meaning that the status of the argument is undecided. Let $\mathtt{in}(\mathcal{L}ab) = \{a \mid \mathcal{L}ab(a) = \mathtt{in}\}$ and $\mathtt{out}(\mathcal{L}ab)$ resp. $\mathtt{undec}(\mathcal{L}ab)$ be defined analogously. The set $\mathtt{in}(\mathcal{L}ab)$ for a labelling $\mathcal{L}ab$ is also called *extension* [Dung, 1995].

We now recall the three basic semantics of argumentation frameworks: the conflict-free, admissible, and complete semantics[2]. Conflict-freeness represents a certain notion of consistency, i. e., we can jointly accept only those arguments that are not in conflict. The intuition behind admissibility is that an argument can only be accepted if all attackers are rejected and if an argument is rejected 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 neither `in` nor `out`.

Definition 2.3. *Let* $\mathcal{L}ab : Ar \to \{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$ *be a labelling for* $AF = (Ar, att)$*. We say that* $\mathcal{L}ab$ *is*

- conflict-free *($\mathcal{CF}$) if for no* $a, b \in \mathtt{in}(L)$ *we have that* $a \in b^-$*,*
- admissible *($\mathcal{AD}$) if and only if it is conflict-free and for every* $a \in Ar$
  - *if* $\mathcal{L}ab(a) = \mathtt{out}$ *then there is* $b \in a^-$ *with* $\mathcal{L}ab(b) = \mathtt{in}$*, and*
  - *if* $\mathcal{L}ab(a) = \mathtt{in}$ *then* $\mathcal{L}ab(b) = \mathtt{out}$ *for all* $b \in a^-$*,*
- complete *($\mathcal{CO}$) if and only if it is admissible and for every* $a \in Ar$*, if* $\mathcal{L}ab(a) = \mathtt{undec}$ *then*
  - *there is no* $b \in a^-$ *s.t.* $\mathcal{L}ab(b) = \mathtt{in}$*, and*
  - *there exists* $c \in a^-$ *s.t.* $\mathcal{L}ab(c) \neq \mathtt{out}$*.*

[2] We observe that in the literature, admissibility and conflict-freeness are sometimes viewed as semantics, and sometimes as properties. Given the increased importance of these notions due to the development of various new kinds of semantics and frameworks, we choose to treat them as semantics in this work.

Different additional types of classical semantics [Dung, 1995; Caminada, 2006; Baroni *et al.*, 2011] can be phrased by imposing further constraints such as minimality or maximality.

Definition 2.4. *Let* $\mathcal{L}ab : Ar \rightarrow \{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$ *be a complete labelling of* $AF = (Ar, att)$*. Then* $\mathcal{L}ab$ *is:*

- grounded *($\mathcal{GR}$) if and only if* $\mathtt{in}(\mathcal{L}ab)$ *is minimal,*
- preferred *($\mathcal{PR}$) if and only if* $\mathtt{in}(\mathcal{L}ab)$ *is maximal, and*
- stable *($\mathcal{ST}$) if and only if* $\mathtt{undec}(\mathcal{L}ab) = \emptyset$*.*

All statements on minimality/maximality are meant to be w.r.t. $\subseteq$*.*

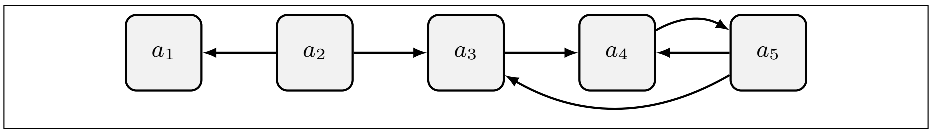

Figure 3: The argumentation framework AF from Example 2.5

Example 2.5. *Consider the abstract argumentation framework AF = (Ar,att) where* $Ar = \{a_1, a_2, a_3, a_4, a_5\}$ *and* $att = \{(a_2, a_1), (a_2, a_3), (a_3, a_4), (a_4, a_5), (a_5, a_4), (a_5, a_3)\}$ *(see also Figure 3). The possible labellings under the admissible, complete, grounded, preferred and stable semantics for this framework are listed in Table 1.*

Labeling	a_1	a_2	a_3	a_4	a_5	$\mathcal{AD}$	$\mathcal{CO}$	$\mathcal{GR}$	$\mathcal{PR}$	$\mathcal{ST}$
$\mathcal{L}ab_1$	undec	undec	undec	undec	undec	✓	×	×	×	×
$\mathcal{L}ab_2$	out	in	out	undec	undec	✓	✓	✓	×	×
$\mathcal{L}ab_3$	undec	undec	out	out	in	✓	×	×	×	×
$\mathcal{L}ab_4$	out	in	out	out	in	✓	✓	×	✓	✓
$\mathcal{L}ab_5$	out	in	out	in	out	✓	✓	×	✓	✓

Table 1: Possible labellings of the framework from Figure 3.

We note that a grounded labelling is uniquely determined and always exists [Dung, 1995] and that not every framework is guaranteed to produce a stable labeling.

One of the most common reasoning problems in argumentation concerns the skeptical and credulous acceptance of arguments, by which we understand that they are accepted in all (resp. some) labellings of a given framework:

Definition 2.6. *Let $\sigma \in \{\mathcal{CF}, \mathcal{AD}, \mathcal{CO}, \mathcal{PR}, \mathcal{GR}, \mathcal{ST}\}$ be any semantics. An argument $a \in Ar$ is* credulously accepted *in AF wrt. σ, written $AF \mid\!\sim^c_\sigma a$, iff $a \in \mathtt{in}(\mathcal{L}ab)$ for some σ-labeling $\mathcal{L}ab$. An argument $a \in Ar$ is* skeptically accepted *in AF wrt. σ, written $AF \mid\!\sim^s_\sigma a$, iff $a \in \mathtt{in}(\mathcal{L}ab)$ for all σ-labelings $\mathcal{L}ab$. We use $\circ \in \{s, c\}$ as a symbol to refer to any inference mode.*

In Example 2.5, argument a_2 would be skeptically accepted under the complete, grounded, preferred and stable semantics. a_4 and a_5 would also be credulously accepted under the complete, preferred and stable semantics, but not under grounded.

We observe that skeptical reasoning is not considered in the case of conflict-free and admissible semantics since a labelling mapping all arguments to `undec` is always conflict–free and admissible. Every argument that does not attack itself will be credulously accepted under the conflict-free semantics.

3 The epistemic approach to probabilistic argumentation

We now go beyond classical three-valued semantics of abstract argumentation and turn to the epistemic approach to probabilistic argumentation [Thimm, 2012; Hunter and Thimm, 2014; Baroni *et al.*, 2014; Hunter and Thimm, 2017]. Instead of evaluating abstract argumentation frameworks by labellings, we rely on probability functions. For this section, we define a probability function as follows.

Definition 3.1. *Let X be some finite set and 2^X its power set. A* probability function *P on X is a function $P : 2^X \to [0, 1]$ that satisfies*

$$\sum_{Y \in 2^X} P(Y) = 1 \quad .$$

Here, a probability function is a function on the set of subsets of some (finite) set which is *normalized*, i.e., the sum of the probabilities of all subsets is one. Let $\mathcal{P}$ be the set of all probability functions.

We use the concept of *subjective probability* [Paris, 1994] for interpreting probabilities. That is, a probability $P(Y)$ for some $Y \subseteq X$ denotes the *degree of belief* we put into Y. Then a probability function P can be seen as an epistemic state of some agent that has uncertain beliefs with respect to X. In probabilistic reasoning [Pearl, 1988; Paris, 1994], this interpretation of probability is widely used to represent and reason over uncertain knowledge.

In the following, we consider probability functions on sets of arguments of an abstract argumentation framework.

Definition 3.2. *Let $AF = (Ar, att)$ be a fixed abstract argumentation framework and let $\mathcal{P}(AF)$ be the set of probability functions of the form $P : 2^{Ar} \to [0,1]$. For $P \in \mathcal{P}(AF)$ and $a \in Ar$, the probability of a is defined as*

$$P(a) = \sum_{a \in Y \subseteq Ar} P(Y) \quad .$$

Given some probability function P, the probability $P(a)$ represents the degree of belief that a is *acceptable* wrt. P. In order to bridge the gap between probability functions and labellings, consider the following definition from [Hunter, 2013].

Definition 3.3. *Let $AF = (Ar, att)$ and $P : 2^{Ar} \to [0,1]$ a probability function on Ar. The labelling $\mathcal{L}ab_P : Ar \to \{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$ defined via the following constraints is called the* epistemic labelling *of P*

- $\mathcal{L}ab_P(a) = \mathtt{in}$ *iff* $P(a) > 0.5$,
- $\mathcal{L}ab_P(a) = \mathtt{out}$ *iff* $P(a) < 0.5$,
- $\mathcal{L}ab_P(a) = \mathtt{undec}$ *iff* $P(a) = 0.5$.

In other words, an argument a is labelled `in` in $\mathcal{L}ab_P$ when it is believed to some degree (which we identify as $P(a) > 0.5$), it is labelled `out` when it is disbelieved to some degree (which we identify as $P(a) < 0.5$), and it is labelled `undec` when it is neither believed nor

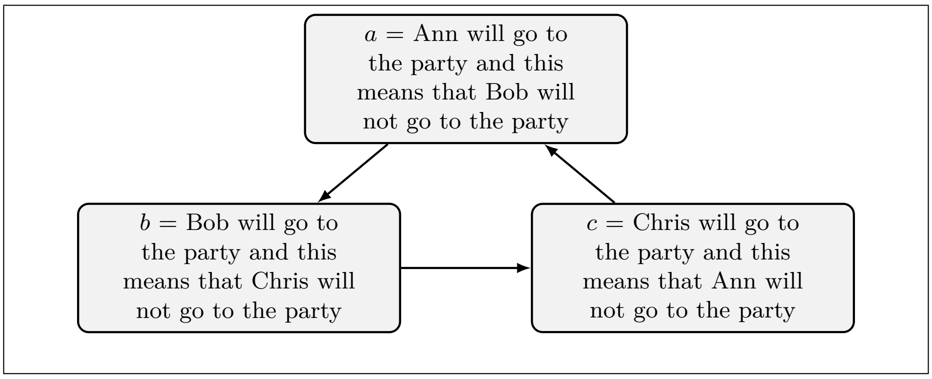

Figure 4: Example of three arguments in a simple cycle.

disbelieved (which we identify as $P(a) = 0.5$). Furthermore, the *epistemic extension* of P is the set of arguments that are labelled `in` by the epistemic labelling, i.e., X is an epistemic extension iff $X = \mathtt{in}(\mathcal{L}ab_P)$. We say that a labelling $\mathcal{L}ab$ and a probability function P are *congruent*, denoted by $\mathcal{L}ab \sim P$, if for all $a \in Ar$ we have $\mathcal{L}ab(a) = \mathtt{in} \Leftrightarrow P(a) = 1$, $\mathcal{L}ab(a) = \mathtt{out} \Leftrightarrow P(a) = 0$, and $\mathcal{L}ab(a) = \mathtt{undec} \Leftrightarrow P(a) = 0.5$. Note that if $\mathcal{L}ab \sim P$ then $\mathcal{L}ab = \mathcal{L}ab_P$, i.e., if a labelling $\mathcal{L}ab$ and a probability function P are congruent then $\mathcal{L}ab$ is also the epistemic labelling of P.

Example 3.4. *To further illustrate epistemic labelings and extensions, consider the graph given in Figure 4. Here, we may believe that, say, a is valid and that b and c are not valid. In which case, with this extra epistemic information about the arguments, we can resolve the conflict and so take the set $\{a\}$ as the "epistemic" extension. In contrast, there is only one admissible set which is the empty set. So by having this extra epistemic information, we get a more informed extension (in the sense that it has harnessed the extra information in a sensible way).*

In general, we want epistemic extensions to allow us to better model the audience of argumentation. In particular, we can deviate from the principles of classical labeling-based semantics for argumentation frameworks. Consider, for example, when a member of the audience of a TV debate listens to the debate at home, she can produce the abstract argument graph based on the arguments and counterarguments exchanged. Then she can identify a probability function to represent the belief she has in each of the arguments. So she may disbelieve some of the arguments based on what she knows about the topic. Furthermore, she may disbelieve some of the arguments that are unattacked. As an extreme, she is at liberty of completely disbelieving all of the arguments (so to assign probability 0 to all of them). If we want to model audiences, where the audience either does not want to or is unable to add counterarguments to an argument graph being constructed in a given argumentation scenario, we need to take the beliefs of the audience into account. Thus, we need to consider which arguments they believe or disbelieve, which may not correspond to how classical argumentation semantics would evaluate the arguments.

Nevertheless, a completely arbitrary probability function is not very informative, and there are various conditions we can impose on probability functions in order to be able to reason with argumentation frameworks in the probabilistic setting. Recall that in the classical case, semantical conditions such as *conflict-freeness* and *admissibility* (see Definition 2.3) play such a role to constrain the set of labellings. In the literature on the epistemic approach, several *rationality postulates* have been proposed to lift these conditions to the probabilistic setting, see also [Hunter and Thimm, 2017] for an overview.

The first property we consider here is a generalization of the *conflict-freeness* property. Using a probability function to interpret an abstract argumentation framework, means that—in order to reason rationally—we cannot have both high degree of belief in an argument and its attacker at the same time. Let $AF = (Ar, att)$ be an abstract argumentation framework and $P : 2^{Ar} \to [0, 1]$ a probability function over Ar. There are two variants how the notion of *conflict-freeness* can be lifted to the probabilistic case:

COH P is *coherent* wrt. AF if for every $a, b \in Ar$, if $(a, b) \in att$ then $P(a) \leq 1 - P(b)$.

RAT P is *rational* wrt. AF if for every $a, b \in Ar$, if $(a, b) \in att$ then $P(a) > 0.5$ implies $P(b) \leq 0.5$.

Both postulates model the general requirement that, if belief in an argument is high, then the belief in an argument attacked by it should be low. While RAT models a rather crisp version of the requirement, COH is a continuous interpretation.

The other important requirement of argumentative reasoning is that, if all attackers of an argument are believed to a rather low belief, then the argument should be believed to a rather high belief. In the classical setting this is the general *reinstatement principle* and implemented through the completeness property [Baroni *et al.*, 2011]. Aspects of this requirement can be modelled through the following postulates:

FOU P is *founded* wrt. AF if $P(a) = 1$ for every $a \in Ar$ with $a^- = \emptyset$.

TRU P is *trusting* wrt. AF if for every a s.t. for every $b \in a^-$, $P(b) < 0.5$, then $P(a) > 0.5$.

OPT P is *optimistic* wrt. AF if $P(a) \geq 1 - \sum_{b \in a^-} P(b)$ for every $a \in Ar$.

The property FOU states that unattacked arguments should receive maximal degree of belief. The property TRU states that arguments whose attackers are disbelieved, should be believed. The property OPT is a generalisation of that idea and states the degree of belief in an argument should be bounded from below by one minus the sum of the beliefs in the attackers. Note that in the special case of having zero belief in all attackers this property requires to have maximal belief in a, just as in the classical case where an argument is accepted if all its attackers are defeated.

A series of further rationality postulates have been proposed in the literature. We refer to [Baroni *et al.*, 2014; Hunter and Thimm, 2017] for a deeper discussion of this topic.

Example 3.5. *Consider the abstract argumentation framework* $AF = (Ar, att)$ *depicted in Figure 5 and the probability functions depicted in Table 2 (note that these functions are only partially defined by giving the probabilities of arguments). The following observations can be made:*

- P_1 *is founded and trusting, but neither rational, coherent, nor optimistic.*

- P_2 *is coherent and rational, but neither founded, trusting nor optimistic.*

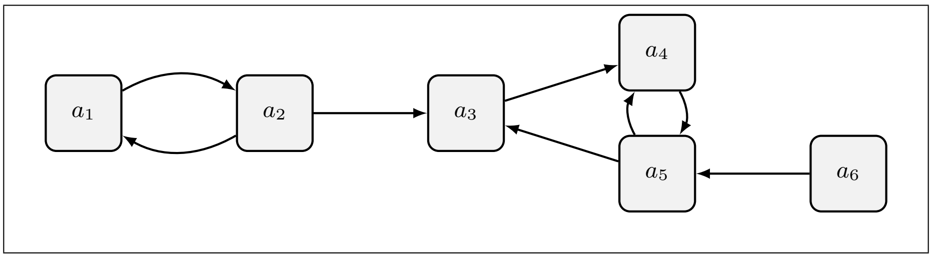

Figure 5: A simple argumentation framework

	a_1	a_2	a_3	a_4	a_5	a_6
P_1	0.2	0.7	0.6	0.3	0.6	1
P_2	0.7	0.3	0.5	0.5	0.2	0.4
P_3	0.7	0.3	0.7	0.3	0	1
P_4	0.7	0.8	0.9	0.8	0.7	1

Table 2: Some probability functions for Example 3.5

- *P_3 is coherent, rational, founded, trusting and optimistic.*
- *P_4 is founded, trusting and optimistic but neither coherent nor rational.*

There are correspondences between probability functions satisfying certain rationality postulates and labellings satisfying certain semantics [Hunter and Thimm, 2017]. For example, if $\mathcal{L}ab$ is an admissible labelling in AF then there is a probability function P satisfying COH and OPT with $\mathcal{L}ab \sim P$ [Thimm, 2012]. Moreover, if $\mathcal{L}ab$ is the grounded labelling and P the one probability function satisfying COH and OPT and having maximal *entropy*[3] then $\mathcal{L}ab \sim P$ [Thimm, 2012].

The framework developed so far already allows for reasoning with abstract argumentation frameworks in a way that incorporates probabilistic interpretations of argumentative principles. For example, fixing a set of rationality postulates, one can determine upper and lower bounds for probabilities of arguments by considering the set of a probability

[3]Define the *entropy* $H(P)$ of P as $H(P) = -\sum_{E \subseteq Ar} P(E) \log P(E)$.

functions satisfying these rationality postulates [Hunter and Thimm, 2017]. If evidence, i. e., bounds or correct degrees of beliefs of some of the arguments, is available, then this can be incorporated in the model and new bounds for the remaining arguments can be calculated [Hunter and Thimm, 2014; Hunter and Thimm, 2017]. In [Hunter and Thimm, 2017], also a method is proposed to allow for probabilistic reasoning if the available evidence is contradictory.

The framework discussed in this section has been extended and analysed in a number of ways. For example, [Gabbay and Rodrigues, 2015] interprets the epistemic approach within the equational approach and [Baroni *et al.*, 2014] extends the epistemic approach to imprecise probabilities. Moreover, Prakken [Prakken, 2017; Prakken, 2018] provides some first thoughts on how probabilistic reasoning can be applied to structured argumentation approaches and relates his ideas with the epistemic approach. In the remainder of this section, we briefly discuss some further extensions.

3.1 Beliefs in Attacks

A natural extension of the epistemic approach is to not only consider degrees of beliefs of arguments but also of attacks [Polberg *et al.*, 2017; Thimm *et al.*, 2018b]. The intuition here is that a high belief in an attack makes the attack *effective* while low belief in attack means that the attack could almost be ignored. This allows for handling both situations in which the belief in an attack is strongly coupled with the belief one has in the argument carrying it out, as well as those where this dependency is much weaker [Polberg and Hunter, 2018].

From the technical side, this generalization can be achieved either by introducing an additional probability function over the attacks of the framework [Polberg *et al.*, 2017] or considering a single probability function of the union of arguments and attacks [Thimm *et al.*, 2018b]. Rationality postulates for the extended setting are then defined by taking uncertainty of attacks into account and new relationships with the classical abstract argumentation can be found [Polberg *et al.*, 2017]. Furthermore, there are also strong relationships with the approach of [Villata *et al.*, 2011] that introduced the notion of acceptability of attacks, see [Thimm *et al.*, 2018b] for a detailed discussion.

3.2 Dynamics

The epistemic approach provides a static perspective on the abstract argumentation framework under consideration. However, as in classical abstract argumentation semantics, we can also consider a more dynamic setting where agents are situated in a particular environment and exchange arguments. There, agents possess opponent models of the other agents and need to update these models as the dialogue progresses. Epistemic probability distributions can be taken as opponent models, causing non-trivial issues to arise when updates have to be performed. In [Hunter and Potyka, 2017; Hunter, 2016] this setting is discussed in more detail and several approaches to updating such models are presented. Section 5.3 of this chapter offers additional discussion on dynamics and probabilistic argumentation.

3.3 Epistemic Graphs

Abstract argumentation frameworks provide a limited expressivity when it comes to the relationships between arguments. The only relationship that can be modelled is the attack, modelling the intuition that an argument cannot be accepted if one of its attackers is accepted. As the epistemic approach is built on top of abstract argumentation it suffers from the issue. However, for abstract argumentation several extensions have been defined to allow for further relationships such as support (see Chapter 1 in this handbook [Cayrol *et al.*, 2021]). To allow for maximal flexibility, *abstract dialectical frameworks* (ADFs) [Brewka *et al.*, 2018; Polberg, 2016] were proposed as a generalisation of abstract argumentation frameworks that allow to define the acceptability function for any argument.

We can find epistemic approaches that aim to combine probabilities with more expressive frameworks. For instance, [Potyka, 2019] introduces several epistemic postulates for modelling support. This can be further generalized to *epistemic graphs* [Hunter *et al.*, 2018b; Hunter *et al.*, 2020; Hunter *et al.*, 2018a; Hunter *et al.*, 2019] which were introduced as a general means for bringing argumentative and probabilistic reasoning capabilities together. Inspired by the freedom of ADFs, epistemic graphs allow the argument graph to be accompanied by a collection of arbitrary probabilistic constraints which impose restrictions between the degrees of belief of different arguments. The existing analysis considers both modelling and technical aspects of epistemic graphs [Hunter *et al.*, 2018b; Hunter *et al.*, 2020; Hunter and Polberg, 2019] as well as their use in a dynamic setting [Hunter *et al.*, 2018a; Hunter *et al.*, 2019], in particular with respect to opponent modelling in dialogues.

4 The constellation approach to probabilistic argumentation

Another natural approach to introduce probabilities into abstract argumentation is to consider a probability distribution over the possible graph structures of the argumentation framework. Every graph structure can then be seen as the true structure in one possible world. In this way, every possible world can again be associated with a set of accepted arguments. Namely, those arguments that are accepted in the particular graph structure under a given classical semantics. The probability of an argument can then again be computed as the probability of being in a world where the argument is accepted. This approach is called the *constellation approach* to probabilistic argumentation. Here we review the constellation approach to probabilistic argumentation from [Hunter, 2012; Hunter and Thimm, 2016], which extends the methods from [Dung and Thang, 2010] and [Li *et al.*, 2011].

4.1 Motivating Examples

To begin with, we illustrate the usefulness of the constellation approach by means of some examples.

Example 4.1. *Consider the following arguments in a legal case:*

- a= *"John says he was not in town when the robbery took place, and therefore denies being involved in the robbery."*
- b= *"Peter says he was at home watching TV when the robbery took place, and therefore denies being involved in the robbery."*

- *c= "Harry says that he is certain that he saw John outside the bank just before the robbery took place, and he also thinks that possibly he saw Peter there too."*

The arguments a and b are arguments that each claim that the speaker is not involved in the robbery, and the argument c is by a potential witness casting doubt on the premises of arguments a and b by undercutting their premises. Also, we see that argument c attacks a with explicit certainty and c attacks b with explicit uncertainty.

If we consider both attacks made by argument c, then we get the argument graph given in Figure 6a. However, if we also take into account the doubt in the attack by c on b, then we get the argument graph given in Figure 6b. This means that there is uncertainty over whether the actual argument graph should be Figure 6a or Figure 6b. We can deal with this uncertainty by regarding the set of spanning subgraphs of Figure 6a (i.e. the four subgraphs given in Figure 6) as a sample space, and assigning a probability to each of them such that the sum is 1. For instance, if Harry has only weak confidence in c attacking b, then the probabilities might be 0.2 for Figure 6a and 0.8 for Figure 6b (i.e. $P(AF_1) = 0.2$ *and* $P(AF_2) = 0.8$*).*

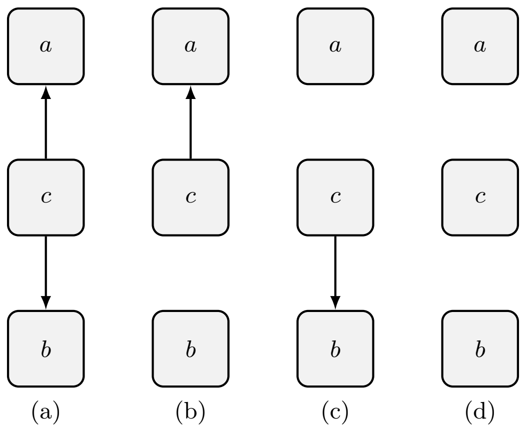

Figure 6: For argument graph AF_1, the subgraphs are (a) AF_1, (b) AF_2, (c) AF_3 and (d) AF_4.

In the example given above, there is explicit uncertainty expressed qualitatively in the attacks made by the arguments. Other situations where uncertainty arises is when there is ambiguity, a form of imprecision in the language used in the arguments, as illustrated in Example 1.1

Another reason for uncertainty in attacks is that real-world arguments presented in natural language are normally enthymemes [Walton, 1989]. An enthymeme is an argument with a support that is insufficient for the claim to be entailed and/or a claim that is incomplete. We see this in Example 1.2 (though Example 1.1 and Example 4.1 also contain enthymemes).

So to summarize, we see at least three kinds of uncertainty that arise in argumentation that we want to capture by quantifying the probability of attack.

Explicit uncertainty of attack Arguments may include some explicit qualification of the attacks made on other arguments. This explicit qualification is usually qualitative (as in Example 4.1), but sometimes it can involve quantitative qualification (such as "I am 99% sure that what John said is a lie").

Implicit imprecision of argument Many arguments have a degree of imprecision in the terminology used (as in Example 1.1). Unless all the language is formally defined, and all participants use the same definitions, it is difficult to avoid some imprecision. This means that when considering two arguments it is not always certain whether or not one attacks the other. For instance, it is possible that "red" and "orange" are consistent together as the description of the same object.

Incomplete premises/claims Most arguments in natural language are enthymemes, which means that they do not explicitly present all their premises and/or claims (as in Example 1.2). With this incompleteness, it is difficult to be certain whether one argument attacks another. If a counterargument has an explicit claim, there may be uncertainty as to whether the attacked argument has the premise that the attacker has contradicted. And if a counterargument has an implicit claim, there may be further uncertainty as to what is being contradicted.

Argumentation often involves multiple agents. This further increases uncertainty in various ways. Consider a typical argumentation scenario where one or more agents are presenting arguments in front of an audience, with the aim of each participant being to persuade the audience to adopt a certain statement. Each participant and the audience have some arguments and counterarguments in mind and may be willing to assimilate further arguments and counterargument.

From an audience's perspective, there may be uncertainty as to what arguments or attacks are in play. The audience may hear various comments in a debate, for example, but they are not sure about the exact set of arguments and attacks that are being put forward. For instance, there may be uncertainty about whether someone has put forward a complex multifaceted argument, or a number of smaller more focused arguments or there may doubt about whether some arguments are just rephrasings of previous arguments. Also, there may be uncertainty about which arguments are meant to be attacked by some argument, which occurs frequently when enthymemes are presented.

From a participant's perspective (i. e., from the perspective of an agent who is about to present arguments and/or attacks during some monological or dialogical argumentation scenario), there may be uncertainty about the audience's opinions, knowledge, values, etc. So when a participant (such as a politician) considers presenting arguments to an audience, the participant might not know for sure which arguments and attacks the audience has in mind. In other words, even before a participant has started, the audience may already have specific arguments and counterarguments in mind and the participant will be adding to those. To handle this, the participant may work with a collection of arguments and counterarguments which he/she assumes will subsume the possibilities for what is held by the audience.

So in general, whether in monological or dialogical situations, we see that there is potentially uncertainty about both arguments and attacks. To address these kinds of uncertainty, we can identify all the possible arguments and attacks that need to be considered. This creates an argument graph, and from this we can identify a probability distribution over the subgraphs of this argument graph. From this distribution, we can then determine the probability that a set of arguments is admissible, or an extension, and the probability that an argument is in an extension.

4.2 Basic definitions

The constellation approach allows us to represent the uncertainty over the topology of the graph. Each subgraph of the original graph is assigned a probability which is understood as the chances of it being the actual argument graph of the agent. It can be used to model what arguments and attacks an agent is aware of. Two important classes of subgraphs are *full subgraphs* that remove arguments, but keep all edges associated with the remaining arguments, and *spanning subgraphs* that keep all arguments, but may remove some of the edges.

Definition 4.2. *Let* $AF = (Ar, att)$ *and* $AF' = (Ar', att')$ *be two argument graphs.* AF' *is a subgraph of* AF*, denoted* $AF' \sqsubseteq AF$*, iff* $Ar' \subseteq Ar$ *and* $att' \subseteq (Ar' \times Ar') \cap att$*.*

- $\wp(AF) = \{AF' \mid AF' \sqsubseteq AF\}$ *denotes the set of all subgraphs of* AF*.*
- *A subgraph* (Ar', att') *is called* **full** *iff* $att' = (Ar' \times Ar') \cap att$*.*
- *A subgraph* (Ar', att') *is called* **spanning** *iff* $Ar' = Ar$ *and* $att' \subseteq att$*.*

Dependent on the application, we may want to restrict our attention to particular subgraphs. If our uncertainty is about which arguments appear in the graph, then only the full (induced) subgraphs of the argument graph have a non–zero probability. If we are only uncertain about which attacks appear, then it is only the spanning subgraphs of the argument graph that can have a non–zero probability.

Definition 4.3. *A* **subgraph distribution** P^c *is a function* $P^c : \wp(AF) \to [0, 1]$ *with* $\sum_{AF' \in \wp(AF)} P^c(AF') = 1$*.*

- P^c *is a* **full subgraph distribution** *iff* $P^c(AF') = 0$ *whenever* AF' *is not a full subgraph.*

- P^c *is a* **spanning subgraph distribution** *iff* $P^c(AF') = 0$ *whenever* AF' *is not a spanning subgraph.*

Note that the above definition follows the notation of [Hunter, 2014], a slightly different notation (but equivalent formalisation) can be found in [Fazzinga *et al.*, 2019].

The constellation approach can be applied for different purposes. One application is to define the probability that a set of arguments or a labeling follows the semantics of a particular type (e. g. grounded, preferred, etc.). This can be done by collecting the probabilities of the subgraphs producing the desired extensions or labelings.

Definition 4.4. *For* $W \subseteq Ar$ *and* $\sigma \in \{\mathcal{CF}, \mathcal{AD}, \mathcal{CO}, \mathcal{PR}, \mathcal{GR}, \mathcal{ST}\}$*, the probability* $P_\sigma(\mathcal{L}ab)$ *that a labeling* $\mathcal{L}ab : W \to \{\mathtt{in}, \mathtt{out}, \mathtt{undec}\}$ *is a* σ*–labeling is defined as:*

$$P_\sigma(\mathcal{L}ab) = \sum_{AF' \in \wp(AF) \ s.t. \ \mathcal{L}ab \in \mathcal{L}_\sigma(AF')} P^c(AF'),$$

where $\mathcal{L}_\sigma(AF')$ *is the set of all* σ*–labelings of* AF'*.*

Another natural application is to define the probability that an argument is accepted under a given semantics.

Definition 4.5. *Given a semantics* $\sigma \in \{\mathcal{AD}, \mathcal{CO}, \mathcal{PR}, \mathcal{GR}, \mathcal{ST}\}$*, the probability that* $a \in Ar$ *is assigned an* `in` *status in a* σ*–labeling is*

$$P_\sigma(a) = \sum_{AF' \in \wp(AF) \ s.t. \ \mathcal{L}ab \in \mathcal{L}_\sigma(AF') \ and \ a \in \mathtt{in}(\mathcal{L}ab)} P^c(AF'),$$

where $\mathcal{L}_\sigma(AF')$ *is the set of all* σ*–labelings of* AF'*.*

We can also define the probability that a set of arguments is an extension.

Definition 4.6. *Given a semantics* $\sigma \in \{\mathcal{AD}, \mathcal{CO}, \mathcal{PR}, \mathcal{GR}, \mathcal{ST}\}$*, the probability* $P_\sigma(W)$ *that a set* $W \subseteq Ar$ *is a* σ*–extension is:*

$$P_\sigma(W) = \sum_{AF' \in \wp(AF) \ s.t. \ \exists \mathcal{L}ab \in \mathcal{L}_\sigma(AF') : W = \mathtt{in}(\mathcal{L}ab)} P^c(AF'),$$

where $\mathcal{L}_\sigma(AF')$ *is the set of all* σ*–labelings of* AF'*.*

Example 4.7. *Consider the graph* $AF = (\{a, b\}, \{(a, b)\})$*. Its subgraphs are* $AF_1 = (\{a, b\}, \{(a, b)\})$*,* $AF_2 = (\{a, b\}, \emptyset)$*,* $AF_3 = (\{a\}, \emptyset)$*,* $AF_4 = (\{b\}, \emptyset)$ *and* $AF_5 = (\emptyset, \emptyset)$ *(see Figure 7). Out of them,* AF_1*,* AF_3*,* AF_4 *and* AF_5 *are full, and* AF_1 *and* AF_2 *are spanning. Consider the following subgraph distribution* P^c*:* $P^c(AF_1) = 0.09$*,* $P^c(AF_2) = 0.81$*,* $P^c(AF_3) = 0.01$ *and* $P^c(AF_4) = 0.09$ *and* $P^c(AF_5) = 0$*. The probability of a given set being a grounded extension is as follows:* $P_{\mathcal{GR}}(\{a, b\}) = P^c(AF_2)$ *= 0.81;* $P_{\mathcal{GR}}(\{a\}) = P^c(AF_1) + P^c(AF_3)$ *= 0.1;* $P_{\mathcal{GR}}(\{b\}) = P^c(AF_4)$ *= 0.09; and* $P_{\mathcal{GR}}(\{\}) = P^c(AF_5)$ *= 0. Therefore* $P_{\mathcal{GR}}(a) = 0.91$ *and* $P_{\mathcal{GR}}(b) = 0.9$*.*

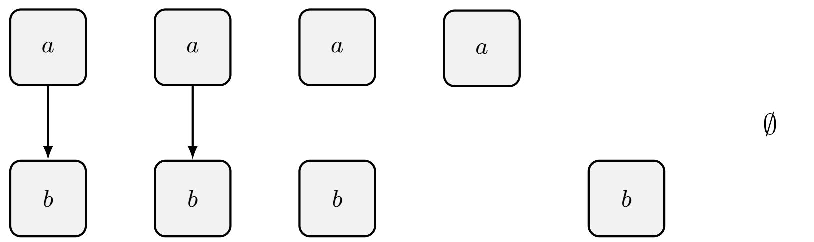

(a) AF (b) AF_1 (c) AF_2 (d) AF_3 (e) AF_4 (f) AF_5

Figure 7: Argument graph AF and its subgraphs.

The proposal for the probabilistic assumption-based argumentation approach of [Dung and Thang, 2010] was the first to consider uncertainty in the topology of the argument graph. It incorporated a probability distribution over sets of arguments, where each set was effectively inducing a subgraph with a probability assignment. It also introduced that the probability of extensions for grounded extensions. Then in a proposal that explicitly considered abstract argumentation, a probability assignment to each argument and to each attack was introduced [Li *et al.*, 2011]. From these two probability functions, a probability distribution over the subgraphs of the argument graph can be obtained. We consider this proposal in detail in the next subsection, and as we will see, it relies on independence between the arguments and between the attacks. This independence assumption does have shortcomings as we will show.

The probability of extensions (first defined by [Dung and Thang, 2010] for grounded semantics) was presented for all semantics in [Li *et* *al.*, 2011]. The probability of labellings was first proposed by [Riveret *et al.*, 2017].

Probabilistic argumentation that is based on explicitly specifying a probability distribution over full subgraphs was first proposed by [Hunter, 2012; Rienstra, 2012]. Subsequently, it was extended to spanning subgraphs [Hunter and Thimm, 2014], and to all subgraphs [Hunter and Thimm, 2016].

For the probabilistic assumption-based argumentation approach of [Dung and Thang, 2010] a dialectical proof procedure was proposed for grounded semantics in [Thang, 2016] and for credulous and ideal semantics in [Hung, 2016b], and a proof procedure based on Bayesian network algorithms was proposed in [Hung, 2016a; Hung, 2017b], as well as algorithms for approximate calculations [Hung, 2018]. In the remainder of this section, we briefly discuss some further extensions.

4.3 Assuming independence

As mentioned above, in [Li *et al.*, 2011] there is a probability assignment to each argument and to each attack. This can be regarded as the uncertainty as to whether the argument or attack should appear in the argument graph. From these two probability functions, a probability distribution over the subgraphs of the argument graph can be obtained. In this subsection, we will review this proposal.

We will consider a simplified version of the proposal by [Li *et al.*, 2011]. In our simplification, we only consider probabilities of arguments (in contrast to the general case of [Li *et al.*, 2011]) where probabilities of attacks are allowed as well).

Definition 4.8. *A* probabilistic argumentation framework PAF *is a triple* $\mathsf{PAF} = (Ar, att, P)$ *where* (Ar, att) *is an abstract argumentation framework and* P *is a function* $P : Ar \to [0, 1]$.

For every argument $a \in Ar$ of a probabilistic argumentation framework PAF the value $P(a)$ is the probability that a is actually present in the argumentation framework. By assuming probabilistic independence between the presence of different arguments, we obtain a probability distribution over sets of arguments. By abuse of notation we denote this probability distribution P as well, which is defined as

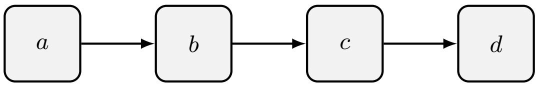

Figure 8: The argumentation framework from Example 4.9

$$P(X) = \prod_{a \in X} P(a) \prod_{a \notin X} (1 - P(a))$$

for all $X \subseteq Ar$. It can be easily shown that $\sum_{X \subseteq Ar} P(X) = 1$, so P is indeed a probability distribution. Given a set $X \subseteq Ar$ of arguments, we denote by $AF{\downarrow_X}$ the induced subgraph of X, i. e. $AF{\downarrow_X} = (X, att \cap (X \times X))$.

Let now $\sigma \in \{\mathcal{CO}, \mathcal{GR}, \mathcal{PR}\}$ be a semantics and $\circ \in \{s, c\}$ be an inference mode. The *probability of acceptance of* a, denoted by $P^{\mathsf{PAF}}_{\circ,\sigma}(a)$, is then defined via

$$P^{\mathsf{PAF}}_{\circ,\sigma}(a) = \sum_{a \in X \subseteq Ar, AF{\downarrow_X} \mid\!\sim^{\circ}_{\sigma} a} P(X) \quad .$$

In other words, $P^{\mathsf{PAF}}_{\circ,\sigma}(a)$ is the sum of the probabilities of the subgraphs of (Ar, att) where a is accepted wrt. σ and $\circ$.

Example 4.9. *Let $AF = (Ar, att)$ be the abstract argumentation frameworks shown in Figure 8 and consider credulous reasoning wrt. grounded semantics. Let $\mathsf{PAF} = (Ar, att, P)$ be a probabilistic argumentation framework with $P(x) = 0.5$ for all $x \in Ar$. Table 3 lists each subset of $X \subseteq Ar$, together with the set of arguments $x \in Ar$ such that $AF{\downarrow_X} \mid\!\sim^{c}_{\mathcal{GR}} x$. For each $X \subseteq Ar$ we have $P(X) = 0.5^4 = 0.0625$. Thus, for each $x \in Ar$ we can calculate the probability $P^{\mathsf{PAF}}_{c,\mathcal{GR}}(x)$ by multiplying the number of subsets of Ar that make x accepted by 0.0625. This yields*

$$P^{\mathsf{PAF}}_{c,\mathcal{GR}}(a) = 0.5, \qquad P^{\mathsf{PAF}}_{c,\mathcal{GR}}(b) = 0.25,$$
$$P^{\mathsf{PAF}}_{c,\mathcal{GR}}(c) = 0.375, \qquad P^{\mathsf{PAF}}_{c,\mathcal{GR}}(d) = 0.3125.$$

Whilst assuming independence brings some advantages, there are situations where it is not appropriate as it does not capture the uncertainty correctly. To illustrate this, we consider the following example.

X	Accepted
$\emptyset$	$\emptyset$
a	a
b	b
a, b	a
c	c
a, c	a, c
b, c	b
a, b, c	a, c
d	d
a, d	a, d
b, d	b, d
a, b, d	a, d
c, d	c
a, c, d	a, c
b, c, d	b, d
a, b, c, d	a, c

Table 3: Choices of X and corresponding accepted arguments in Example 8

Example 4.10. *Consider the argument graph AF_1 in Figure 9 where the meaning for the arguments is as follows.*

- a = *According to John, Peter hit John first.*
- b = *According to Peter, John hit Peter first.*

Assuming that John and Peter are two children having a playground contratemps, and each is accusing the other of having thrown the first punch, we may regard the argument graph AF_1 as a good reflection of the situation. However, if we also suppose that Peter is regularly getting into fights with other children in the playground, then we might want to use probabilistic argumentation. For instance, we may regard AF_2 as the most likely scenario, AF_3 as less likely but not impossible, and AF_6 as impossible because there has clearly been punching and someone is lying. This can easily be captured in the constellations approach with

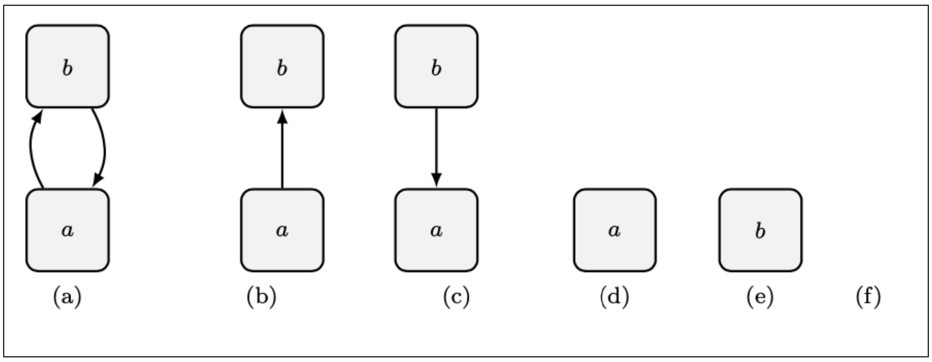

Figure 9: For argument graph AF_1, the subgraphs are (a) AF_1, (b) AF_2, (c) AF_3, (d) AF_4, (e) AF_5, and (f) AF_6 (which is the empty graph).

a probability distribution over the graphs where for instance, argument graph AF_2 has probability 0.9 and argument graph AF_3 has probability 0.1. But using Definition 4.8, we see we cannot obtain this distribution. We can only get a probability distribution over AF_1, AF_4, AF_5, and AF_6. This is because the definition forces us to drop arguments and the attacks that involve them, but does not allow us to just drop attacks.

As we said earlier, in this review, we only provided a simplified version of the proposal by [Li *et al.*, 2011]. In our simplification, we only consider probabilities of arguments (in contrast to the general case of [Li *et al.*, 2011]) where probabilities of attacks are allowed as well). Using the original version, we still would be unable to get the probability distribution of AF_2 having probability 0.9 and argument graph $AF3$ having probability 0.1. The following example, is a further illustration of the shortcoming of the independence assumption over attacks.

Example 4.11. *Consider argument graph AF_1 in Figure 10a where the meaning for the arguments is as follows.*

- a = *CheapAir is going bust.*
- b = *The CheapAir tickets to Paris are a bargain; We should buy them for our holiday.*
- c = *The CheapAir tickets to New York are a bargain; We should buy them for our conference trip.*

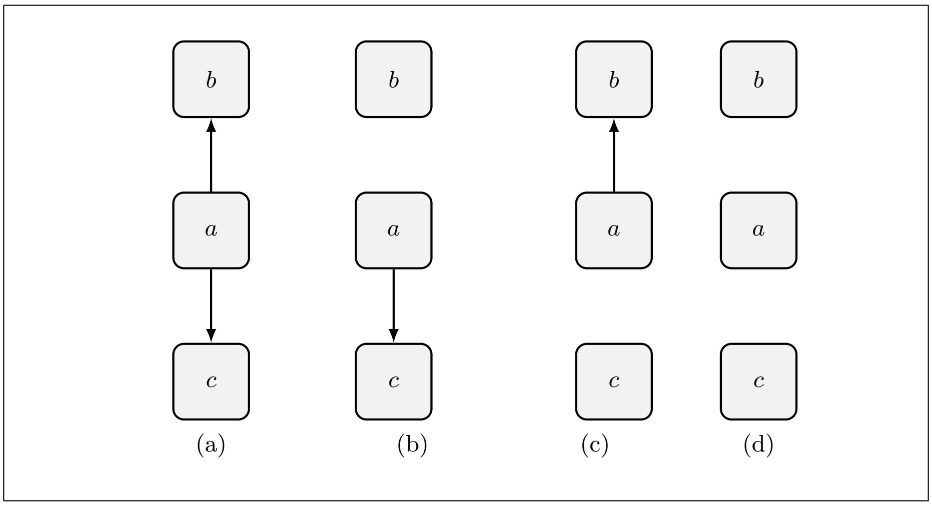

Figure 10: For argument graph AF_1, the subgraphs are (a) AF_1, (b) AF_2, (c) AF_3 and (d) AF_4.

Here the attack by a on b and the attack by a on c are not independent. For instance, if the attack by a on b is shown to be true, then there is a raised probability that the attack by a on c is true.

So the above example illustrates how assuming independence over attacks means that the example is not correctly modelled. In contrast, if we assume a probability distribution over the subgraphs, then the example can be correctly modelled as illustrated in the following example.

Example 4.12. *Continuing Example 4.11, if we start with graph AF_1 in Figure 10a, then there are four subgraphs to consider. Suppose we let the probability function over subgraphs be the following:*

$$P(AF_1) = 0.7, \qquad P(AF_2) = 0, \qquad P(AF_3) = 0, \qquad P(AF_4) = 0.3.$$

The marginals for the attacks are as follows:

$$\begin{aligned} P(a,b) &= P(AF_1) + P(AF_3) = 0.7, \\ P(a,c) &= P(AF_1) + P(AF_2) = 0.7. \end{aligned}$$

So in some settings, it is reasonable to start with a probability distribution over arguments (and attacks), and assume independence between the arguments (and attacks) in order to construct a probability distribution over the subgraphs. In general, it may not be justified to assume independence.

4.4 Further work

For further reading concerning the constellation approach to probabilistic argumentation we refer the readers to [Hunter, 2012; Dung and Thang, 2010; Li *et al.*, 2011; Hunter, 2013; Dondio, 2014a; Hunter, 2014]. Computational results can be found in [Dondio, 2014a; Dondio, 2014b; Fazzinga *et al.*, 2013; Fazzinga *et al.*, 2015; Fazzinga *et al.*, 2019]. In [Doder and Woltran, 2014], we can find a characterization of one of the versions of the constellation approach in terms of probabilistic logic.

Whilst using the constellations approach is computationally expensive [Fazzinga *et al.*, 2015], developments in approximation [Fazzinga *et al.*, 2013; Fazzinga *et al.*, 2018a; Fazzinga *et al.*, 2016b; Fazzinga *et al.*, 2016a; Alfano *et al.*, 2020] and automated reasoning could be harnessed [Bistarelli *et al.*, 2018; Mantadelis and Bistarelli, 2020].

The constellations approach has been generalized to bipolar argumentation [Fazzinga *et al.*, 2018c; Fazzinga *et al.*, 2018b], and to abstract dialectical frameworks [Polberg and Doder, 2014], and it has been used as the basis of a proposal for graded semantics [Thimm *et al.*, 2018a]. In order to deal with uncertainty in the probability assignment, the constellations approach has been extended to support upper and lower bounds on the probability assignments by using credal sets [Morveli-Espinoza *et al.*, 2019], and extended to support Dempster-Shafer theory [Hung, 2017a; Hung, 2017c].

Applications of the constellations approach include modelling the argument graph held by opponents in argumentative discussions [Hunter and Thimm, 2016], aggregating different perspectives on an argument graph [Hunter and Noor, 2020], and evaluating plans [Hung *et al.*, 2015].

5 Other Topics in Probabilistic Argumentation

In the previous sections, we reviewed the two main approaches to probabilistic argumentation, namely the epistemic approach and the constellation approach. In this section, we will give an overview of some other topics in probabilistic argumentation. To begin with, we review some ideas for combining the epistemic and the constellation approach. Subsequently, we will look at some ideas to apply probabilistic techniques in order to learn argumentation frameworks from data. Probabilistic ideas have also been applied to model changing beliefs in argumentation. We will review some ideas in the third subsection. Finally, we sketch some ideas about how argumentation technology can be applied to enhance other probabilistic models.

5.1 Combinations of the Epistemic and the Constellations approach

The *Labelling Framework for Probabilistic Argumentation* [Riveret *et al.*, 2018] is a framework that subsumes and combines a number of approaches considered before. The framework, based on a defeasible logic instantiation of abstract argumentation, deals with three representations of uncertainty, where each representation induces the next. These are probabilistic theory frames (PTFs), probabilistic graph frames (PGFs), and probabilistic labelling frames (PLFs). Roughly speaking, a PTF is a probability distribution over subsets of rules of a given knowledge base. A PTF induces a PGF, which is a probability distribution over subgraphs consisting of arguments constructed on the basis of a set of rules. An argument is either *on* (included) or *off* (not included) in a subgraph. Given a labelling-based semantics, a PGF induces a PLF, which is a probability distribution over labellings. Finally, a PLF can be used to calculate the probability that a statement is accepted. The notion of PGF subsumes the constellations approach but does not enforce independence of different arguments being on or off. Furthermore it is shown that there exists a correspondence between the notion of PLF and a probability distribution over extensions as used in the epistemic approach described in Section 3. Some of the postulates for the epistemic approach are shown to hold under various assumptions about the labelling-based semantics that is used. The presented framework represents a combination of both the constellations and epistemic approach, as it deals with both uncertainty about the topology of an argumentation framework, which arises from uncertainty about elements of the knowledge base, as well as the resulting uncertainty about whether argument is accepted. A similar combination of the constellations and epistemic approach appears in [Rienstra, 2012].

5.2 Learning approaches

The work [Riveret and Governatori, 2016] proposed an anytime algorithm for learning the structure of argumentation graphs from a data stream of labellings. The authors restrict their discussion to grounded labellings, but many ideas can be extended easily to other semantics. The basic idea is to start from a complete graph over the arguments and to collect information about the plausibility of attack edges as labellings arrive. For example, if two arguments are labelled in simultaneously, the attack edge between them can be eliminated. The authors present an algorithm based on 4 rules that allow increasing or decreasing the credibility of edges online while receiving labelling information. If the labellings in the stream are sampled from an unknown argumentation graph such that every labelling has a positive probability of being sampled, then the algorithm will eventually find an argumentation graph that is equivalent to the original graph (in the sense that they have the same complete extensions) with probability 1.

In a similar spirit, [Kido and Okamoto, 2017] consider the problem of estimating the attack relation between arguments from the acceptance statuses given by different users. That is, given sets of arguments accepted by different users, find an attack relation that explains the accepted arguments. To this end, the authors propose a Bayesian network model that defines the probability of extensions and acceptability statues based on the attack relation and the acceptability semantics that can be preferred, stable, grounded or complete. The authors propose definitions for the local probability models in their Bayesian network. Their main assumptions are that larger attack relations are less likely, that all semantics are equally likely and that all extensions under a given semantics and attack relation are equally likely. Probabilistic inference methods can be used to compute the most likely attack relation under the Bayesian network model given the sets of accepted arguments and an acceptability semantics. The authors demonstrate the applicability in an online discussion scenario and prove some analytical guarantees. In general, the true attack relation may not be identifiable from data alone, but in some special cases, it is.

[Riveret *et al.*, 2017] train Boltzman machines on sets of grounded labellings and demonstrate that the trained models can be used to generate grounded labellings with a similar distribution. The sampling algorithm is sound in the sense that it is guaranteed to return only grounded labellings and complete in the sense that every grounded labelling can be generated.

5.3 Dynamics of argumentation and dialogues

The work [Hunter, 2015] proposed epistemic probabilistic argumentation as a model for the belief state of agents in persuasion dialogues. Roughly speaking, the current belief state of the persuadee is represented as a probability distribution over extensions. The persuader can posit arguments that the persuadee may accept or reject. Dependent on the persuadee's judgement, the current belief state can be updated. For this purpose, different update functions have been proposed in [Hunter, 2015] that basically redistribute probability mass such that the persuadee's feedback is taken into account. These update functions can be employed to simulate the outcome of possible persuasion dialogues in order to find an optimal persuasion strategy. Another family of update operators based on the idea of minimizing the distance to the prior beliefs has been introduced in [Hunter and Potyka, 2017] and compared to the redistribution update operators. The language for updates has been extended in [Hunter *et al.*, 2018a] and a fragment that guarantees polynomial runtime guarantees can be found in [Potyka *et al.*, 2019].

In [Shakarian *et al.*, 2016], Belief Revision has been considered for structured probabilistic argumentation. The approach combines Nilsson's probabilistic logic [Nilsson, 1986] and defeasible logic programming with presumptions [Martinez *et al.*, 2012]. Probabilistic logic is used to define an *environmental model* that captures case-specific knowledge which may be uncertain. Defeasible logic programming with presumptions is used to define an *analytical model* that contains general background knowledge which may be inconsistent when combined naively. Both models are combined via an *annotation function* that connects formulas in the environmental and analytical model. The need for belief revision results from the fact that probabilistic logic requires consistent assumptions even if defeasible logic programming does not. The authors study different types of inconsistencies and ways to perform belief revision on the environmental model, the analytical model and the annotation function. In particular, postulates for revising the analytical model and the annotation function are proposed similar to the more general belief revision literature [Hansson, 1997; Falappa *et al.*, 2012] and representation theorems are provided. The usefulness of the approach is demonstrated by an apllication in cyber security [Shakarian *et al.*, 2016]. The approach has been refined in [Simari *et al.*, 2016] by taking quantitative aspects into account like how the probabilities change due to revision.

5.4 Using argumentation for probabilistic models

While this survey deals mainly with probabilistic approaches in modelling uncertainty in argumentation, there are also approaches which use argumentation as a tool to reason about probabilistic models.

[Timmer *et al.*, 2017] proposed an argumentation-based method for explaining inferences in Bayesian networks. They note that Bayesian networks, while widely used and well-understood, are hard to interpret for non-statistical experts. Thus, argumentation may help in explaining Bayesian networks in a way that corresponds to everyday reasoning. Starting with a Bayesian network together with a chosen variable of interest, their method is based on constructing a *support graph*, which consists of reasoning chains that start with potential pieces of evidence and end with the variable of interest. The support graph is then used to construct arguments which are represented within the ASPIC+ argumentation framework [Modgil and Prakken, 2014]. However, the use of argumentation in this setting is somewhat limited when compared with typical argumentation-based methods, as one argument attacks another only if the latter is weaker. As a result, the grounded extension of the resulting argumentation framework coincides with the set of arguments that receive no attacks. This grounded extension acts as an intuitive explanation for the calculation of the likelihood ratio or posterior probability of the variable of interest.

The work [Nielsen and Parsons, 2007] developed an argumentation-based method for Bayesian network fusion. The setting they consider is of a set of agents where each agent is equipped with a Bayesian network representing the agent's domain model. They then address the problem of having the agents compromise and agree on a single Bayesian network. Their model builds on an extension of Dung's [1995] abstract argumentation model supporting set-based attacks. An argumentation framework is defined such that the preferred extensions correspond to possible compromises, as well as a formal multi-agent debate model where a debate ends with finding one of these compromises. The debate model enjoys the property that a debate is guaranteed to end up with the best possible compromise according to a given compromise score function. Note that the method developed here is concerned purely with obtaining a compromise on the *structure* of the agent's Bayesian networks. It does not specify how probabilistic parameterisations of the different Bayesian networks are be fused.

6 Summary and Conclusion

The two main approaches to probabilistic (abstract) argumentation are the constellations and the epistemic approaches. In the constellations approach, there is uncertainty over the topology of the argument graph, whereas in the epistemic approach, the topology of the argument graph is fixed, but there is uncertainty about whether an argument is believed. The epistemic approach has been extended to also allow a probability distribution over subsets of attacks, and thereby represent belief in each attack. A further approach is based on labellings for arguments using in, out, and undecided, augmented with off for denoting that the argument does not occur in the graph. A probability distribution over labellings gives a form of probabilistic argumentation that overlaps with the constellations and epistemic approaches.

There are many other ideas of how to combine probability theory and argumentation, some of which we sketched in the previous section. Others, in particular those relating to probabilistic reasoning with structured argumentation such as [Haenni, 2009] had to left out of this survey.

One particular natural application is learning argumentation graphs from data. Since data is usually noisy and uncertain, probability theory is a natural tool in this area. The additional expressiveness of probabilistic argumentation also raises new questions when it comes to dynamics of argumentation and novel ideas have been studied to model changing beliefs in or by means of probabilistic models. Conversely, argumentation technology has also been used to enhance other probabilistic models, for example, by adding explanation capabilities.

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 *et al.*, 2010]. Argumentation can also be used to help construct Bayesian networks [Bex and Renooij, 2016; Wieten *et al.*, 2019]. 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]. 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 [Nielsen and Parsons, 2007].

Looking forward, we envisage some important developments that will harness key advantages of probabilistic argumentation. These include explaining knowledge learned from data using probabilistic argumentation (see for example [Hunter, 2020]), aggregating and analysing potentially conflicting information from multiple sources (see for example [Hunter and Noor, 2020]), using probabilistic argumentation to dynamically model participants during dialogues so that strategically good moves are made (see for example [Hunter, 2015; Hunter and Thimm, 2016; Hunter *et al.*, 2019]), and using probabilistic argumentation for modelling non-normative behaviour by participants (see for example [Polberg and Hunter, 2018]).

Acknowledgments

The work reported in this chapter has been partially supported by the Deutsche Forschungsgemeinschaft (grants 375588274 and 432308570).

References

[Alfano *et al.*, 2020] G. Alfano, M. Calautti, S. Greco, F. Parisi, and I. Trubitsyna. Explainable acceptance in probabilistic abstract argumentation: Complexity and approximation. In *Proceedings of the 17th International Conference on Principles of Knowledge Representation and Reasoning*, pages 33–43, 2020.

[Baroni *et al.*, 2011] P. Baroni, M. Caminada, and M. Giacomin. An introduction to argumentation semantics. *The Knowledge Engineering Review*, 26(4):365–410, 2011.

[Baroni *et al.*, 2014] P. Baroni, M. Giacomin, and P. Vicig. On rationality conditions for epistemic probabilities in abstract argumentation. In Simon Parsons, Nir Oren, Chris Reed, and Federico Cerutti, editors, *Proceedings of COMMA'14*, volume 266 of *FAIA*, pages 121–132. IOS Press, 2014.

[Besnard and Hunter, 2008] P. Besnard and A. Hunter. *Elements of Argumentation*. The MIT Press, 2008.

[Bex and Renooij, 2016] 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.

[Bistarelli *et al.*, 2018] S. Bistarelli, T. Mantadelis, F. Santini, and C. Taticchi. Probabilistic argumentation frameworks with metaproblog and conarg. In Lefteri H. Tsoukalas, Éric Grégoire, and Miltiadis Alamaniotis, editors, *Proceedings of ICTAI'18*, pages 675–679. IEEE, 2018.

[Brewka *et al.*, 2014] G. Brewka, S. Polberg, and S. Woltran. Generalizations of Dung frameworks and their role in formal argumentation. *IEEE Intelligent Systems*, 29(1):30–38, Jan 2014.

[Brewka *et al.*, 2018] G. Brewka, S. Ellmauthaler, H. Strass, J. P. Wallner, and S. Woltran. Abstract dialectical frameworks. In Pietro Baroni, Dov Gabbay, Massimiliano Giacomin, and Leendert van der Torre, editors, *Handbook of Formal Argumentation*, chapter 5. College Publications, February 2018.

[Caminada, 2006] M. Caminada. Semi-stable semantics. In *Proceedings of the First International Conference on Computational Models of Argumentation (COMMA'06)*, pages 121–130, 2006.

[Cayrol *et al.*, 2021] Claudette Cayrol, Andrea Cohen, and Marie-Christine Lagasquie-Schiex. Higher-order interactions (bipolar or not) in abstract argumentation: A state of the art. In Dov Gabbay, Massimiliano Giacomin, Guillermo R. Simari, and Matthias Thimm, editors, *Handbook of Formal Argumentation*, volume 2, chapter 1. College Publications, 2021.

[Doder and Woltran, 2014] D. Doder and S. Woltran. Probabilistic argumentation frameworks – a logical approach. In Umberto Straccia and Andrea

Calì, editors, *Proceedings of SUM'14*, volume 8720 of *LNCS*, pages 134–147. Springer, 2014.

[Dondio, 2014a] P. Dondio. Multi-valued and probabilistic argumentation frameworks. In *Proceedings of the Fifth International Conference on Computational Models of Argumentation (COMMA'14)*, 2014.

[Dondio, 2014b] P. Dondio. Toward a computational analysis of probabilistic argumentation frameworks. *Cybernetics and Systems*, 45(3):254–278, 2014.

[Dung and Thang, 2010] P. M. Dung and P. M. Thang. Towards (probabilistic) argumentation for jury-based dispute resolution. In Pietro Baroni, Federico Cerutti, Massimiliano Giacomin, and Guillermo Ricardo Simari, editors, *Proceedings of COMMA'10*, volume 216 of *FAIA*, pages 171–182. IOS Press, 2010.

[Dung, 1995] 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.

[Falappa *et al.*, 2012] M. A Falappa, G. Kern-Isberner, . DL Reis, and G. R Simari. Prioritized and non-prioritized multiple change on belief bases. *Journal of Philosophical Logic*, 41(1):77–113, 2012.

[Fazzinga *et al.*, 2013] B. Fazzinga, S. Flesca, and F. Parisi. On the complexity of probabilistic abstract argumentation. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI'13)*, 2013.

[Fazzinga *et al.*, 2015] B. Fazzinga, S. Flesca, and F. Parisi. On the complexity of probabilistic abstract argumentation frameworks. *ACM Transactions on Computational Logic*, 16(3), May 2015.

[Fazzinga *et al.*, 2016a] B. Fazzinga, S. Flesca, and F. Parisi. On efficiently estimating the probability of extensions in abstract argumentation frameworks. *International Journal of Approximate Reasoning*, 69:106–132, 2016.

[Fazzinga *et al.*, 2016b] B. Fazzinga, S. Flesca, F. Parisi, and A. Pietramala. Computing or estimating extensions' probabilities over structured probabilistic argumentation frameworks. *IfCoLog Journal of Logics and their Applications*, 3(2):177–200, 2016.

[Fazzinga *et al.*, 2018a] B. Fazzinga, S. Flesca, and F. Furfaro. Credulous and skeptical acceptability in probabilistic abstract argumentation: complexity results. *Intelligenza Artificiale*, 12(2):181–191, 2018.

[Fazzinga *et al.*, 2018b] B. Fazzinga, S. Flesca, and F. Furfaro. Probabilistic bipolar abstract argumentation frameworks: complexity results. In Jérôme Lang, editor, *Proceedings of IJCAI'18*, pages 1803–1809. International Joint Conferences on Artificial Intelligence, 2018.

[Fazzinga *et al.*, 2018c] B. Fazzinga, S. Flesca, F. Furfaro, and F. Scala. Computing extensions' probabilities over probabilistic bipolar abstract argumentation frameworks. In Pierpaolo Dondio and Luca Longo, editors, *Proceedings of AI^3@AI*IA'18*, volume 2296 of *CEUR Workshop Proceedings*, pages 57–70. CEUR-WS.org, 2018.

[Fazzinga *et al.*, 2019] B. Fazzinga, S. Flesca, and F. Furfaro. Complexity of fundamental problems in probabilistic abstract argumentation: Beyond independence. *Artificial Intelligence*, 268:1–29, 2019.

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

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

[Grabmair *et al.*, 2010] 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)*, pages 255–266, 2010.

[Haenni, 2009] R. Haenni. Probabilistic argumentation. *Journal of Applied Logic*, 7(2):155 – 176, 2009. Special issue: Combining Probability and Logic.

[Hansson, 1997] S. Hansson. Semi-revision. *Journal of Applied Non-Classical Logics*, 7(1-2):151–175, 1997.

[Hung *et al.*, 2015] N. Duy Hung, S. Kaisaard, and S. Hnin Pwint Oo. Probabilistic argumentation for service restoration in power distribution networks. In T. Theeramunkong, A. M. J. Skulimowski, T. Yuizono, and S. Kunifuji, editors, *Recent Advances and Future Prospects in Knowledge, Information and Creativity Support Systems - Selected Revised Papers from the Tenth International Conference on Knowledge, Information and Creativity Support Systems (KICSS 2015), 12-14 November 2015, Phuket, Thailand*, volume 685 of *Advances in Intelligent Systems and Computing*, pages 54–68. Springer, 2015.

[Hung, 2016a] N. Duy Hung. Computing probabilistic assumption-based argumentation. In Richard Booth and Min-Ling Zhang, editors, *Proceedings of PRICAI'16*, volume 9810 of *LNCS*, pages 152–166. Springer, 2016.

[Hung, 2016b] N. Duy Hung. A probabilistic argumentation engine. In Nikolaos Bourbakis, Anna Esposito, Amol Mali, and Miltos Alamaniotis, editors, *Proceedings of ICTAI'16*, pages 311–318. IEEE, 2016.

[Hung, 2017a] N. Duy Hung. A generalization of probabilistic argumentation with dempster-shafer theory. In Gabriele Kern-Isberner, Johannes Fürnkranz, and Matthias Thimm, editors, *Proceedings of KI'17*, volume 10505 of *LNCS*, pages 155–169. Springer, 2017.

[Hung, 2017b] N. Duy Hung. Inference procedures and engine for probabilistic argumentation. *International Journal of Approximate Reasoning*, 90:163–191, 2017.

[Hung, 2017c] N. Duy Hung. Probabilistic assumption-based argumentation with DST evidence. In *Proceedings of IFSA-SCIS'17*, pages 1–6. IEEE, 2017.

[Hung, 2018] N. Duy Hung. Progressive inference algorithms for probabilistic argumentation. In Tim Miller, Nir Oren, Yuko Sakurai, Itsuki Noda, Bastin Tony Roy Savarimuthu, and Tran Cao Son, editors, *Proceedings of PRIMA'18*, volume 11224 of *LNCS*, pages 371–386. Springer, 2018.

[Hunter and Noor, 2020] A. Hunter and K. Noor. Aggregation of perspectives using the constellations approach to probabilistic argumentation. In *Proceedings of AAAI'20*, 2020.

[Hunter and Polberg, 2019] A. Hunter and S. Polberg. A model-based theorem prover for epistemic graphs for argumentation. In Gabriele Kern-Isberner and Zoran Ognjanović, editors, *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, pages 50–61, Cham, 2019. Springer International Publishing.

[Hunter and Potyka, 2017] A. Hunter and N. Potyka. Updating probabilistic epistemic states in persuasion dialogues. In Alessandro Antonucci, Laurence Cholvy, and Odile Papini, editors, *Proceedings of ECSQARU'17*, volume 10369 of *LNCS*, pages 46–56. Springer, 2017.

[Hunter and Thimm, 2014] A. Hunter and M. Thimm. Probabilistic argumentation with incomplete information. In Torsten Schaub, Gerhard Friedrich, and Barry O'Sullivan, editors, *Proceedings of ECAI'14*, volume 263 of *FAIA*, pages 1033–1034. IOS Press, 2014.

[Hunter and Thimm, 2016] A. Hunter and M. Thimm. Optimization of dialectical outcomes in dialogical argumentation. *International Journal of Approximate Reasoning*, 78:73–102, 2016.

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

[Hunter *et al.*, 2018a] A. Hunter, S. Polberg, and N. Potyka. Updating Belief in Arguments in Epistemic Graphs. In Michael Thielscher, Francesca Toni, and Frank Wolter, editors, *Proceedings of KR'18*, pages 138–147. AAAI Press, 2018.

[Hunter *et al.*, 2018b] A. Hunter, S. Polberg, and M. Thimm. Epistemic graphs for representing and reasoning with positive and negative influences of arguments. *ArXiv*, 2018.

[Hunter *et al.*, 2019] A. Hunter, S. Polberg, and N. Potyka. Delegated updates in epistemic graphs for opponent modelling. *International Journal of Approximate Reasoning*, 113:207 – 244, 2019.

[Hunter *et al.*, 2020] A. Hunter, S. Polberg, and M. Thimm. Epistemic graphs for representing and reasoning with positive and negative influences of arguments. *Artificial Intelligence*, 281:103236, 2020.

[Hunter, 2012] A. Hunter. Some foundations for probabilistic abstract argumentation. In Bart Verheij, Stefan Szeider, and Stefan Woltran, editors, *Proceedings of COMMA'12*, volume 245 of *FAIA*, pages 117–128. IOS Press, 2012.

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

[Hunter, 2014] A. Hunter. Probabilistic qualification of attack in abstract argumentation. *International Journal of Approximate Reasoning*, 55:607–638, 2014.

[Hunter, 2015] A. Hunter. Modelling the persuadee in asymmetric argumentation dialogues for persuasion. In *Proceedings of IJCAI'15*, 2015.

[Hunter, 2016] A. Hunter. Persuasion dialogues via restricted interfaces using probabilistic argumentation. In Steven Schockaert and Pierre Senellart, editors, *Proceedings of SUM'16*, volume 9858 of *LNCS*, pages 184–198. Springer, 2016.

[Hunter, 2020] A. Hunter. Generating instantiated argument graphs from probabilistic information. In *Proceedings of ECAI'20*, 2020.

[Kido and Okamoto, 2017] H. Kido and K. Okamoto. A bayesian approach to argument-based reasoning for attack estimation. In Carles Sierra, editor, *Proceedings of IJCAI'17*, pages 249–255. International Joint Conferences on Artificial Intelligence, 2017.

[Li *et al.*, 2011] H. Li, N. Oren, and T.J. Norman. Probabilistic argumentation frameworks. In Sanjay Modgil, Nir Oren, and Francesca Toni, editors, *Proceedings of TAFA'11*, volume 7132 of *LNCS*, pages 1–16. Springer, 2011.

[Mantadelis and Bistarelli, 2020] T. Mantadelis and S. Bistarelli. Probabilistic abstract argumentation frameworks, a possible world view. *Int. J. Approx. Reason.*, 119:204–219, 2020.

[Martinez *et al.*, 2012] M. V. Martinez, A. J. García, and G. R. Simari. On the use of presumptions in structured defeasible reasoning. In Bart Verheij, Stefan Szeider, and Stefan Woltran, editors, *Computational Models of Argument - Proceedings of COMMA 2012*, volume 245 of *Frontiers in Artificial Intelligence and Applications*, pages 185–196. IOS Press, 2012.

[Modgil and Prakken, 2014] S. Modgil and H. Prakken. The ASPIC+ framework for structured argumentation: a tutorial. *Argument and Computation*, 5:31–62, 2014.

[Morveli-Espinoza *et al.*, 2019] M. Morveli-Espinoza, J. C. Nieves, and C. A. Tacla. An imprecise probability approach for abstract argumentation based on credal sets. In Gabriele Kern-Isberner and Zoran Ognjanović, editors, *Proceedings of ECSQARU'19*, volume 11726 of *LNCS*, pages 39–49. Springer, 2019.

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

[Nilsson, 1986] N. J Nilsson. Probabilistic logic. *Artificial intelligence*, 28(1):71–87, 1986.

[Paris, 1994] J. B. Paris. *The Uncertain Reasoner's Companion – A Mathematical Perspective.* Cambridge University Press, 1994.

[Pearl, 1988] J. Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference.* Springer, 1988.

[Polberg and Doder, 2014] S. Polberg and D. Doder. Probabilistic abstract dialectical frameworks. In Eduardo Fermé and João Leite, editors, *Proceedings of JELIA'14*, volume 8761 of *LNCS*, pages 591–599. Springer, 2014.

[Polberg and Hunter, 2018] 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.

[Polberg *et al.*, 2017] S. Polberg, A. Hunter, and M. Thimm. Belief in attacks in epistemic probabilistic argumentation. In Serafín Moral, Olivier Pivert, Daniel Sánchez, and Nicolás Marín, editors, *Proceedings of SUM'17*, volume 10564 of *LNCS*, pages 223–236. Springer, 2017.

[Polberg, 2016] S. Polberg. Understanding the Abstract Dialectical Framework. In Loizos Michael and Antonis Kakas, editors, *Proceedings of JELIA'16*, volume 10021 of *LNCS*, pages 430–446. Springer, 2016.

[Pollock, 1995] J. L. Pollock. *Cognitive Carpentry: A Blueprint for How to Build a Person.* The MIT Press, Cambridge, Massachusetts, 1995.

[Potyka *et al.*, 2019] N. Potyka, S. Polberg, and A. Hunter. Polynomial-time updates of epistemic states in a fragment of probabilistic epistemic argumentation. In *Proceedings of ECSQARU'19*, 2019.

[Potyka, 2019] N. Potyka. A polynomial-time fragment of epistemic probabilistic argumentation. *International Journal of Approximate Reasoning*, 115:265 – 289, 2019.

[Prakken, 2017] H. Prakken. On relating abstract and structured probabilistic argumentation: A case study. In Alessandro Antonucci, Laurence Cholvy, and Odile Papini, editors, *Proceedings of ECSQARU'17*, volume 10369 of *LNCS*, pages 69–79. Springer, 2017.

[Prakken, 2018] H. Prakken. Probabilistic strength of arguments with structure. In Michael Thielscher, Francesca Toni, and Frank Wolter, editors, *Proceedings of KR'18*, pages 158–167. AAAI Press, 2018.

[Rienstra, 2012] T. Rienstra. Towards a probabilistic dung-style argumentation system. In *Proceedings of AT'12*, 2012.

[Riveret and Governatori, 2016] R. Riveret and G. Governatori. On learning attacks in probabilistic abstract argumentation. In Catholijn M. Jonker, Stacy Marsella, John Thangarajah, and Karl Tuyls, editors, *Proceedings of AAMAS'16*, pages 653–661. IFAAMAS, 2016.

[Riveret *et al.*, 2017] R. Riveret, D. Korkinof, M. Draief, and J. Pitt. Probabilistic abstract argumentation: An investigation with boltzmann machines. *Argument & Computation*, 8(1):89, 2017.

[Riveret *et al.*, 2018] R. Riveret, P. Baroni, Y. Gao, G. Governatori, A. Rotolo, and G. Sartor. A labelling framework for probabilistic argumentation. *Annals of Mathematics and Artificial Intelligence*, 83(1):21–71, 2018.

[Shakarian *et al.*, 2016] P. Shakarian, G.I. Simari, G. Moores, D. Paulo, S. Parsons, M. A Falappa, and A. Aleali. Belief revision in structured probabilistic argumentation. *Annals of Mathematics and Artificial Intelligence*, 78(3-4):259–301, 2016.

[Simari *et al.*, 2016] G. I. Simari, P. Shakarian, and M. A. Falappa. A quantitative approach to belief revision in structured probabilistic argumentation. *Annals of Mathematics and Artificial Intelligence*, 76(3-4):375–408, 2016.

[Thang, 2016] P. M. Thang. Dialectical proof procedures for probabilistic abstract argumentation. In Matteo Baldoni, Amit K. Chopra, Tran Cao Son, Katsutoshi Hirayama, and Paolo Torroni, editors, *Proceedings of PRIMA'16*, volume 9862 of *LNCS*, pages 397–406. Springer, 2016.

[Thimm *et al.*, 2018a] M. Thimm, F. Cerutti, and T. Rienstra. Probabilistic graded semantics. In Sanjay Modgil, Katarzyna Budzynska, and John Lawrence, editors, *Proceedings of COMMA'18*, volume 305 of *FAIA*, pages 369–380. IOS Press, 2018.

[Thimm *et al.*, 2018b] M. Thimm, S. Polberg, and A. Hunter. Epistemic attack semantics. In Sanjay Modgil, Katarzyna Budzynska, and John Lawrence, editors, *Proceedings of the Seventh International Conference on Computational Models of Argumentation (COMMA'18)*, volume 305 of *Frontiers in Artificial Intelligence and Applications*, pages 37–48, Warsaw, Poland, September 2018.

[Thimm, 2012] M. Thimm. A probabilistic semantics for abstract argumentation. In Luc De Raedt, Christian Bessiere, Didier Dubois, Patrick Doherty, Paolo Frasconi, Fredrik Heintz, and Peter Lucas, editors, *Proceedings of ECAI'12*, volume 242 of *FAIA*, pages 750–755. IOS Press, 2012.

[Timmer *et al.*, 2015] 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.

[Timmer *et al.*, 2017] S.T. Timmer, J.Ch. Meyer, H. Prakken, S. Renooij, and B. Verheij. A two-phase method for extracting explanatory arguments from bayesian networks. *International Journal of Approximate Reasoning*, 80:475 – 494, 2017.

[Toni, 2014] F. Toni. A tutorial on assumption-based argumentation. *Argument & Computation*, 5(1):89–117, 2014.

[Verheij, 2014] B. Verheij. Arguments and their strength: Revisiting pollock's anti-probabilistic starting points. In Simon Parsons, Nir Oren, Chris Reed, and Federico Cerutti, editors, *Proceedings of COMMA'14*, volume 266 of *FAIA*, pages 433–444. IOS Press, 2014.

[Villata *et al.*, 2011] S. Villata, G. Boella, and L. van der Torre. Attack Semantics for Abstract Argumentation. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence (IJCAI'11)*, pages 406–413, 2011.

[Vreeswijk, 2004] G. A. W. Vreeswijk. Argumentation in bayesian belief networks. In Iyad Rahwan, Pavlos Moraïtis, and Chris Reed, editors, *Proceedings of ArgMAS'04*, volume 3366 of *LNCS*, pages 111–129. Springer, 2004.

[Walton, 1989] D. Walton. *Informal Logic: A Handbook for Critical Argumentation*. Cambridge University Press, 1989.

[Wieten *et al.*, 2019] R. Wieten, F. Bex, H. Prakken, and S. Renooij. Supporting discussions about forensic bayesian networks using argumentation. In *Proceedings of ICAIL'19*, pages 143–152, 2019.

[Wu and Caminada, 2010] Y. Wu and M. Caminada. A labelling-based justification status of arguments. *Studies in Logic*, 3(4):12–29, 2010.