

Aggregation of Perspectives Using the Constellations Approach to Probabilistic Argumentation

Anthony Hunter, Kawsar Noor

Department of Computer Science
University College London
Gower Street, London, WC1E 6BT, UK
{anthony.hunter, kawsar.noor.15}@ucl.ac.uk

Abstract

In the constellations approach to probabilistic argumentation, there is a probability distribution over the subgraphs of an argument graph, and this can be used to represent the uncertainty in the structure of the argument graph. In this paper, we consider how we can construct this probability distribution from data. We provide a language for data based on perspectives (opinions) on the structure of the graph, and we introduce a framework (based on general properties and some specific proposals) for aggregating these perspectives, and as a result obtaining a probability distribution that best reflects these perspectives. This can be used in applications such as summarizing collections of online reviews and combining conflicting reports.

Introduction

Uncertainty can arise in argumentation for various reasons. For an individual argument, there can be uncertainty in the premises of the argument, or in whether the claim follows from the premises. For instance, many arguments are enthymemes, which means that an argument might not explicitly present all its premises and/or claim. So when an argument is heard or read by someone, that recipient has to decode the enthymeme to recover the intended argument. There is therefore uncertainty in this decoding process as the decoded argument could be different to the intended argument. So using enthymemes introduces uncertainty as to the content of individual arguments, and uncertainty as to whether one argument attacks another argument.

In a study on dialogical argumentation, participants were asked to categorize relationships between arguments as one of attack, support, dependent somehow, or unrelated, and the results showed that whilst there was a good agreement between participants on the categorization for the majority of pairs of arguments, there was a sizeable subset for which there was some disagreement (Polberg and Hunter 2018).

Further uncertainty arises in dialogical argumentation. When one agent is presenting an argument to another agent, the agent presenting the argument is unsure what related arguments the intended recipient is aware of, and what beliefs the intended recipient might have concerning those arguments. For the agent presenting the argument, this might be

important if, for example, the agent wants to persuade the other agent to accept a specific argument. In this kind of scenario, the persuader has to make good choices of argument based on what arguments and attacks she thinks the persuadee is aware of (Hunter and Thimm 2016).

Probabilistic approaches for modeling uncertainty in argumentation include the constellations approach and the epistemic approach (Hunter 2013). The first is based on a probability distribution over the subgraphs of the argument graph ((Hunter 2012) which extends (Dung and Thang 2010) and (Li, Oren, and Norman 2011)), and this can be used to represent the uncertainty over the topology of the graph (i.e. whether a particular argument or attack appears in the argument graph). The second is the epistemic approach which can be used to represent the degree to which each argument is believed (Thimm 2012; Hunter 2013; Hunter and Thimm 2017). A further approach is based on labellings for arguments using *in*, *out*, and *undecided*, from (Caminada and Gabbay 2009), augmented with *off* for denoting that the argument does not occur in the graph (Riveret and Governatori 2016). A probability distribution over labellings gives a form of probabilistic argumentation that overlaps with the constellations and epistemic approaches.

In this paper, we focus on the constellations approach, and consider how we can automatically construct a probability distribution over a constellation of argument graphs. For this, we will use perspectives (i.e. opinions) on what the probability distribution should be. Consider for example a number of reviews on a product. Each review contains some arguments that taken together indicate the overall sentiment that the writer has towards the product. Obviously, an argument graph is not given by the writer, but the reader may consider an implicit argument graph that includes the writer's arguments. So a question is then what is the implicit argument graph? Since there is unlikely to be a single argument graph that would reflect all the reviews, we can harness the constellations approach to represent the uncertainty over the topology. Our approach is to provide a language for individual perspectives (represented by constraints on the extensions of the argument graph), and then provide methods for taking a tuple of perspectives and obtaining a probability distribution over a constellation of argument graphs.

To illustrate, we can consider reports on whether a particular drug treatment is good. There are websites such as




Figure 1: A constellation of argument graphs where the arguments are: (A) “Drug is good”; (B) “Drug is not good”; (C) “There are side-effects”; and (D) “There are notable benefits”. In this simple example, there is uncertainty about whether there are attacks by C on A and by D on B .

webmd.com and askapatient.com that allow a patient to review a drug in terms of the benefits and side-effects of the treatment. These can be regarded as arguments that taken together provide an assessment of the treatment. So the arguments extracted from a review give a perspective. We give a very simple example of a constellation of argument graphs for drug reviews in Figure 1. We use a very simple graph in this paper for reasons of space. But for modelling a set of real reviews, we can consider more arguments such as for specific types of benefit and side-effect as well as further counterarguments to individual benefits (e.g. benefit effect wears off too quickly) or side-effect (e.g. nausea side-effect can be treated with anti-emetics). So often, the graph contains a more complex structure than just pros and cons.

Potential applications of the proposal in this paper include analyzing conflicting reports, e.g. online reviews (Noor, Hunter, and Mayer 2017), sense-making in intelligence gathering, e.g. (Cerutti et al. 2018), understanding the opinions of individuals prior to a policy decision, e.g. (Krauthoff et al. 2018). and user modeling from crowdsourced data that can be used for strategic construction of persuasion dialogues (Hunter and Thimm 2016).

Preliminaries

An **argument graph** G is a pair $(\mathcal{A}, \mathcal{R})$ where \mathcal{A} is a set and $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ (Dung 1995). Let $\text{Nodes}(G) = \mathcal{A}$ and $\text{Arcs}(G) = \mathcal{R}$.

Each element $\alpha \in \mathcal{A}$ is called an **argument** and $(\alpha, \beta) \in \mathcal{R}$ means that α **attacks** β (accordingly, α is said to be an **attacker** or a **counterargument** for β). A set of arguments $S \subseteq \mathcal{A}$ **attacks** $\beta \in \mathcal{A}$ iff there is an argument $\alpha \in S$ such that α attacks β . Also, S **defends** $\alpha' \in \mathcal{A}$ iff for each argument $\beta \in \mathcal{A}$, if β attacks α' then S attacks β . A set $S \subseteq \mathcal{A}$ of arguments is **conflict-free** iff there are no arguments $\alpha, \alpha' \in S$ such that α attacks α' . Let Γ be a conflict-free set of arguments, and let $\text{Defended} : \wp(\mathcal{A}) \rightarrow \wp(\mathcal{A})$ be a function such that $\text{Defended}(\Gamma) = \{\alpha \mid \Gamma \text{ defends } \alpha\}$.

Dialectical semantics are given to argument graphs by *extensions*, i.e., sets of arguments that are considered to be jointly acceptable. We consider the following types of extensions: (i) Γ is a **complete extension** (co) iff $\Gamma = \text{Defended}(\Gamma)$, (ii) Γ is a **grounded extension** (gr) iff it is the (uniquely determined) minimal (w.r.t. set inclusion) complete extension, (iii) Γ is a **preferred extension** (pr) iff it is a maximal (w.r.t. set inclusion) complete extension, and (iv) Γ is a **stable extension** (st) iff it is a preferred exten-

sion such that Γ attacks β for each argument $\beta \in \Gamma \setminus \mathcal{A}$. For $G = (\mathcal{A}, \mathcal{R})$, let $\text{Extensions}_\sigma(G)$ be the set of extensions of G according to semantics $\sigma \in \{\text{co}, \text{pr}, \text{gr}, \text{st}\}$.

Given a set of arguments \mathcal{A} , we define a language \mathcal{L} to represent and reason with any argument graph G where $\text{Nodes}(G) \subseteq \mathcal{A}$. This language is a sublanguage of the *logic for dialectical outcomes* (Hunter and Thimm 2016) for representing and reasoning about extensions of a graph.

Definition 1 (Syntax for language). *Let \mathcal{A} be a set of arguments. Let \mathcal{L} be the modal language defined as follows: $\phi \in \mathcal{L}$ iff it is defined via the BNF: $\phi ::= \psi \mid \neg\phi \mid \phi \vee \phi \mid \phi \wedge \phi \mid \diamond\phi$, where $\psi \in \mathcal{A}$.*

The language \mathcal{L} is based on each atom representing an argument and Boolean combinations of arguments using negation, disjunction, and conjunction, plus an \diamond operator. The formulae are used to describe extensions.

Definition 2 (Semantics for language). *For a graph G , let $E \subseteq \text{Nodes}(G) \subseteq \mathcal{A}$, $\psi, \phi \in \mathcal{L}$. The **satisfaction relation** for a dialectical semantics σ , denoted \models_σ , is defined as follows.*

$$\begin{aligned} G, E \models_\sigma \alpha & \text{ iff } \alpha \in E \\ G, E \models_\sigma \diamond\phi & \text{ iff for some } E' \in \text{Extensions}_\sigma(G), \\ & \text{ we have } G, E' \models_\sigma \phi \\ G, E \models_\sigma \phi \wedge \psi & \text{ iff } G, E \models_\sigma \phi \text{ and } G, E \models_\sigma \psi \\ G, E \models_\sigma \phi \vee \psi & \text{ iff } G, E \models_\sigma \phi \text{ or } G, E \models_\sigma \psi \\ G, E \models_\sigma \neg\phi & \text{ iff } G, E \not\models_\sigma \phi \end{aligned}$$

Furthermore define $G \models_\sigma \phi$ iff for all $E \in \text{Extensions}_\sigma(G)$ we have $G, E \models_\sigma \phi$.

Example 1. *Consider Figure 1. Some inferences include the following: $G_1 \models_{\text{gr}} \neg A \wedge \neg B \wedge C \wedge D$; $G_2 \models_{\text{gr}} A \wedge C \wedge D$; and $G_4 \models_{\text{pr}} \diamond A \wedge \diamond B$.*

We also need to introduce some notions for subgraphs of an argument graph. Let $\mathcal{R} \otimes \mathcal{A}'$ be the subset of \mathcal{R} involving just the arguments in $\mathcal{A}' \subseteq \mathcal{A}$, i.e., $\mathcal{R} \otimes \mathcal{A}' = \{(\alpha, \beta) \in \mathcal{R} \mid \alpha, \beta \in \mathcal{A}'\}$. Also let G_\emptyset denote the empty graph. For argument graphs $G = (\mathcal{A}, \mathcal{R})$ and $G' = (\mathcal{A}', \mathcal{R}')$ we say that G' is a **subgraph** of G , denoted $G' \sqsubseteq G$, iff $\mathcal{A}' \subseteq \mathcal{A}$ and $\mathcal{R}' \subseteq \mathcal{R} \otimes \mathcal{A}'$. For any argument graph G , let $\text{Subgraphs}(G)$ denote the set of subgraphs of G (i.e. $\{G' \mid G' \sqsubseteq G\}$).

Let G^* be the **complete graph** for \mathcal{A} (i.e. $\text{Nodes}(G^*) = \mathcal{A}$ and $\text{Arcs}(G^*) = \{(A, B) \mid A, B \in \mathcal{A}\}$). A set of graphs is a **constellation set** for G^* , denoted \mathcal{G} , if for all $G \in \mathcal{G}$, $G \sqsubseteq G^*$. We assume that the correct subgraph is in \mathcal{G} , and so in Definition 3 all graphs not in \mathcal{G} have zero probability. If we are uncertain about all the arguments and attacks, then $\mathcal{G} = \text{Subgraphs}(G^*)$, whereas if we are uncertain about which arguments (respectively attacks) are in the correct subgraph, then \mathcal{G} is the set of full (respectively spanning) subgraphs of G^* . As another example, we could assume a set of necessary arcs \mathcal{R}^n , and a set of possible arcs \mathcal{R}^p , so that $\mathcal{G} = \{(\mathcal{A}, \mathcal{R}^n \cup \mathcal{R}') \mid \mathcal{R}' \subseteq \mathcal{R}^p\}$, and so the uncertainty is about which of the possible arcs is in the actual graph. To, illustrate, consider Figure 1, where $\mathcal{G} = \{G_1, G_2, G_3, G_4\}$, the necessary arcs are $\{(A, B), (B, A)\}$ and the optional arcs are $\{(C, A), (D, B)\}$.

Definition 3. Let \mathcal{G} be a constellation set. A **probability distribution** for \mathcal{G} is a function $P : \mathcal{G} \rightarrow [0, 1]$ such that $\sum_{G \in \mathcal{G}} P(G) = 1$. Let \mathcal{P} denote the set of probability distributions for \mathcal{G} .

The acceptability probability of a formula is the sum of the probability of the graphs in \mathcal{G} that imply it.

Definition 4. Let \mathcal{G} be a constellation set and let $\sigma \in \{\text{co, pr, gr, st}\}$ be a semantics. Also let P be a probability distribution for \mathcal{G} . For a $\phi \in \mathcal{L}$, the **probability of acceptability** of ϕ w.r.t. σ is $P_\sigma(\phi) = \sum_{G \in \mathcal{G} \text{ s.t. } G \models_\sigma \phi} P(G)$.

Example 2. Consider Figure 1, with $P(G_1) = 0.5$, $P(G_2) = 0.2$, $P(G_3) = 0.2$, and $P(G_4) = 0.1$. Some inferences that we can consider include the following: $P_{\text{gr}}(C \wedge D) = P(G_1) + P(G_2) + P(G_3) + P(G_4) = 1$; $P_{\text{gr}}(B \wedge C \wedge D) = P(G_3) + P(G_4) = 0.3$; and $P_{\text{pr}}(\diamond A \wedge \diamond B) = P(G_4) = 0.1$.

The notion of probability of acceptability subsumes the definition of the probability that a set of arguments is an extension, as well as the definition for the probability that an argument is an inference, cf. (Li, Oren, and Norman 2011; Hunter 2012; 2013). For every argument graph G and semantics $\sigma \in \{\text{co, pr, gr, st}\}$, (Hunter and Thimm 2016) show that: If for all graphs G , $G \models_\sigma \phi$, implies $G \models_\sigma \psi$, then $P_\sigma(\phi) \leq P_\sigma(\psi)$; And if $\phi \wedge \psi$ is unsatisfiable, then $P_\sigma(\phi \vee \psi) = P_\sigma(\phi) + P_\sigma(\psi)$.

For certain applications a restricted set of probability distributions can be used where the probability values come from a finite set of values (Hunter, Polberg, and Thimm 2018). This may be appropriate if we want to represent probability values as in a Likert scale (Likert 1931). It also has the benefit of always producing a finite set of answers. However, for the approach to be coherent, this set should meet certain basic requirements. We thus use the notion of a reasonable restricted value set, which has to be closed under addition and subtraction (assuming the resulting value is in the $[0, 1]$ interval) and contains 1.

Definition 5. A finite set of rational numbers from the unit interval Π is a **reasonable restricted value set** iff $1 \in \Pi$ and for any $x, y \in \Pi$ it holds that if $x + y \leq 1$, then $x + y \in \Pi$, and if $x - y \geq 0$, then $x - y \in \Pi$.

Since, we will only consider reasonable restricted value sets, we will refer to them as restricted value sets. Examples include $\{0, 1\}$, $\{0, 0.5, 1\}$, and $\{0, 0.25, 0.5, 0.75, 1\}$.

Perspectives

In the functions we define in the rest of this paper, we assume that they are with respect to a restricted value set Π , a set of arguments \mathcal{A} , a constellation set \mathcal{G} , and a semantics σ . Whilst these should be given as indices on the functions, we will omit them to ease the notation.

Next, we consider how agents can give their views on a probability distribution over \mathcal{G} .

Definition 6. A **profile element** is a pair (ϕ, v) where ϕ is a formula in \mathcal{L} , and v is in a restricted value set Π . Let \mathcal{X} denote the set of all profile elements that can be formed from \mathcal{L} and Π . A **profile** is a tuple $[x_1, \dots, x_n]$ where each x_i is

Graph Structure	G_1 $A \leftrightarrow B$	G_2 $A \rightarrow B$	G_3 $A \leftarrow B$	G_4 $A \quad B$
P_1	1	0	0	0
P_2	0	1	0	0
P_3	0	0	1	0
P_4	0	0	0	1
P_5	0.5	0.5	0	0
P_6	0.5	0	0.5	0
P_7	0.5	0	0	0.5
P_8	0	0.5	0.5	0
P_9	0	0.5	0	0.5
P_0	0	0	0.5	0.5

Table 1: If $\Pi = \{0, 1\}$, (resp. $\Pi = \{0, 0.5, 1\}$), then $\{P_1, \dots, P_4\}$ (resp. $\{P_1, \dots, P_0\}$) are the distributions.

a profile element. Let $\mathcal{T} = \{[x_1, \dots, x_n] \mid x_1, \dots, x_n \in \mathcal{X}\}$ be the set of profiles.

Each profile element comes from an agent (e.g. via a review by the agent). We may choose to restrict each agent to only giving one profile element, but this is not necessary. A profile therefore collects the different views on the probability distribution.

Definition 7. A probability distribution P satisfies a profile element $x = (\phi, v)$ wrt a semantics σ iff $P_\sigma(\phi) = v$. Let $\text{Sat}(x)$ be the set of probability distributions that satisfy (ϕ, v) wrt a semantics σ . For the empty profile $T = []$, $\text{Sat}(T) = \text{Sat}(\top)$, and for a profile $T = [x_1, \dots, x_n]$, $\text{Sat}(T) = \text{Sat}(x_1) \cap \dots \cap \text{Sat}(x_n)$.

The least (respectively most) constraining profile element is (\top, v) (respectively (\perp, v)) for any $v \in \Pi$ since any (respectively no) probability distribution will satisfy it.

Example 3. For $\sigma = \text{gr}$, and $\Pi = \{0, 1\}$ (i.e. P_1, \dots, P_4 in Table 1), $\text{Sat}((A \wedge \neg B, 1)) = \{P_2\}$, and $\text{Sat}((\neg A \wedge B, 1)) = \{P_3\}$. But, $\text{Sat}([(A \wedge \neg B, 1), (\neg A \wedge B, 1)]) = \emptyset$.

Returning to our motivating example in Figure 1, we can use a profile element to represent how a reviewer might view the argument graph. For instance, the element $(A \wedge \neg B, 1)$ denotes that it is certain that A and not B is in the extension, the element $(A \vee B, 1)$ denotes that it is certain that either A or B is in the extension (and so it is completely vague as to the extension), and the element $(\diamond A, 0)$ denotes that it is certain that A is not in a preferred extension.

We now provide some subsidiary definitions. For profile elements $x = (\phi, v)$ and $x' = (\phi', v')$, x is **equivalent** to x' iff $\vdash \phi \leftrightarrow \phi'$ and $v = v'$. For profiles T and T' , T is equivalent to T' (denoted $T \equiv T'$) iff there is a permutation of $T = [x_1, \dots, x_n]$ and $T' = [x'_1, \dots, x'_n]$ such that for each i , x_i is equivalent to x'_i . Finally, $T + T'$ denotes the **concatenation** of profiles T and T' .

Next, we define the Form function which takes a probability distribution and returns a profile that represents it. Since, for each profile, there are potentially many equivalent profiles, we assume a representative for each set of equivalent profiles though how this is done is not important for this paper.




Figure 2: Example of argument graphs.

Definition 8. Let \mathcal{T} be a set of profiles and let \mathcal{P} be a set of probability distributions. The function $\text{Form} : \mathcal{P} \rightarrow \mathcal{T}$ is defined as follows: $\text{Form}(P) = [(\phi_1, v_1), \dots, (\phi_n, v_n)]$ where the following three conditions hold for $i, j \in \{1, \dots, n\}$: (1) $\phi_i \not\vdash \phi_j$, (2) $v_i \neq 0$; and (3) for each $v \in (0, 1]$, $P(\psi) = v$ iff $\sum_{i \in \{1, \dots, n\}} \text{s.t. } \phi_i \vdash \psi v_i = v$.

We explain the conditions as follows: (1) there is no redundancy of elements; (2) all elements are non-zero; and (3) the probability of a formula is the sum of the probability of the elements that entail it.

Example 4. Consider Table 1. For P_8 , and $\sigma = \text{gr}$, $\text{Form}(P_8) = [(A \wedge \neg B, 0.5), (\neg A \wedge B, 0.5)]$. And for P_1 , and $\sigma = \text{pr}$, $\text{Form}(P_1) = [(\Diamond A \wedge \Diamond B, 1)]$.

Definition 9. A profile T is maximal iff there is a probability distribution P such that $\text{Form}(P) = T$.

Proposition 1. If T is maximal, and $P \in \text{Sat}(T)$, then $\text{Form}(P) \equiv T$.

Example 5. Consider Table 1 and $\sigma = \text{gr}$. $T = [(A, 1)]$ is not maximal since $\text{Sat}(T) = \{P_2, P_4, P_9\}$, and for all $P_i \in \{P_2, P_4, P_9\}$, $\text{Form}(P_i) \neq T$.

Essentially, a maximal profile encodes the probability of extensions of the subgraphs for a probability distribution.

Example 6. For Table 1, if P is P_2 , and $\sigma = \text{gr}$, then $P_\sigma(A \wedge \neg B) = 1$, and so the maximal profile is $T = [(A \wedge \neg B, 1)]$.

Proposition 2. Let P be a probability distribution, and let T be a maximal profile such that $\text{Form}(P) = T$. There is an element (ϕ, v) in T iff there are subgraphs $G_1, \dots, G_k \in \mathcal{G}$ s.t. for each $i \in \{1, \dots, k\}$, $E \in \text{Extensions}_\sigma(G_i)$ and $\vdash \bigwedge E \wedge \bigwedge_{\alpha \in A \setminus E} \neg \alpha \leftrightarrow \phi$ and $\sum_{i \in \{1, \dots, n\}} P(G_i) = v$.

Because more than one graph in a constellation set can have the same extension, we can have a profile being satisfied by multiple probability distributions.

Proposition 3. For a probability distribution P , $P \in \text{Sat}(\text{Form}(P))$.

Example 7. Consider Figure 2 where \mathcal{G} is $\{G_1, G_2\}$ and $\Pi = \{0, 1\}$ and $\mathcal{P} = \{P_1, P_2\}$ s.t. $P_1(G_1) = 1$ and $P_2(G_2) = 1$. For each graph, the extension is $\{A\}$ and so $\text{Form}(P_1) = \text{Form}(P_2) = (A \wedge \neg B \wedge \neg C, 1)$. Also $\text{Sat}(A \wedge \neg B \wedge \neg C, 1) = \{P_1, P_2\}$, and $\text{Form}(\text{Sat}(A \wedge \neg B \wedge \neg C, 1)) = (A \wedge \neg B \wedge \neg C, 1)$. Therefore, $P_1 \in \text{Sat}(\text{Form}(P_1))$ and $P_2 \in \text{Sat}(\text{Form}(P_2))$.

As we use profiles to represent different perspectives (opinions) on a probability distribution, we need principled ways to identify that distribution, and in particular when the profile T is such that the cardinality of $\text{Sat}(T)$ is greater than 1, or the profile is inconsistent (i.e. when $\text{Sat}(T) = \emptyset$). We

address this need using an induction function that takes a profile and returns a probability distribution that is the best compromise for the views in the profile. We define this in terms of aggregation and combination as follows.

Aggregation

Given a profile, an aggregation function finds the set of probability distributions that best reflect the profile.

Definition 10. Let \mathcal{T} be the set of profiles and \mathcal{P} be the set of probability distributions. An **aggregation function** is a function $F : \mathcal{T} \rightarrow \wp(\mathcal{P})$.

We will focus on aggregation functions that satisfy the following general properties.

- **Viability:** $F(T) \neq \emptyset$
- **Consistency:** If $\text{Sat}(T) \neq \emptyset$, then $F(T) = \text{Sat}(T)$.
- **Equivalence:** If $T_1 \equiv T_2$, then $F(T_1) = F(T_2)$.
- **Fairness:** If $\text{Sat}([x_1]) \cap \text{Sat}([x_2]) = \emptyset$, then $F([x_1, x_2]) \not\subseteq \text{Sat}(x_2)$.
- **Concordance:** $F([x_1]) \cap F([x_2]) \subseteq F([x_1, x_2])$

We explain these general properties as ensuring that: (Viability) there is always a set of distributions as the aggregation of a profile; (Consistency) the set of distributions for the aggregations are those of the profile when the profile is consistent; (Equivalence) the set of distributions for two profiles are the same when the profiles are equivalent; (Fairness) when two elements are contradictory, then neither element has its satisfying distributions in the aggregation; and (Concordance) if the aggregation for each of two profiles agree on a distribution, then that distribution is in the aggregation of the two profiles. These properties are inspired by postulates for belief merging operators that take a tuple of knowledge bases and return a single knowledge base (Konieczny and Pérez 2002), whereas our properties take a tuple of profile elements (formula and belief in that formula) and return a probability distribution.

The following result shows commutativity (i.e. the order of elements in the profiles does not affect the aggregation, and it is direct consequence of equivalence).

Proposition 4. For $T_1, T_2 \in \mathcal{T}$, if T_1 is a permutation of T_2 , then $F(T_1) = F(T_2)$.

The following property of monotonicity states that adding profile elements to the profile will result in a subset of probability distributions being returned.

- **Monotonicity:** For $T_1, T_2 \in \mathcal{T}$, if $T_1 \subseteq T_2$, then $F(T_2) \subseteq F(T_1)$.

Monotonicity may appear to be a natural property for aggregation but it cannot hold with the properties that we will assume for an aggregation function.

Proposition 5. If F satisfies consistency and fairness, then it cannot satisfy monotonicity.

Though if we restrict consideration to consistent profiles, then we do get monotonicity as a direct consequence of the consistency property.

Proposition 6. Let F satisfy consistency. For $T_1, T_2 \in \mathcal{T}$, if $T_1 \subseteq T_2$, and $\text{Sat}(T_2) \neq \emptyset$, then $F(T_2) \subseteq F(T_1)$.

We discriminate between two types of aggregation function (again inspired by (Konieczny and Pérez 2002)) using the following properties where \mathbb{N} is the set of natural numbers, $k \in \mathbb{N}$, and $[x]^k$ denotes a profile $[x_1, \dots, x_k]$ s.t. $x = x_1 = \dots = x_k$.

- **Majority:** For each profile T , and probability distribution P , there is a profile element x and a $k \in \mathbb{N}$ s.t. $F(T + [x]^k) = \{P\}$.
- **Arbitration:** For each profile T , profile element x , and $k \in \mathbb{N}$, $F(T + [x]^k) = F(T + [x])$.

An aggregation function that satisfies majority aims to minimize dissatisfaction of the profile as a whole, whereas an aggregation function that satisfies arbitration aims to minimize dissatisfaction of any individual profile element.

Proposition 7. If an aggregation function satisfies fairness, then it cannot satisfy both majority and arbitration.

We now consider some specific aggregation functions that we define in terms of a ranking over probability distributions using the following measure, where \mathbb{R} is the set of reals.

Definition 11. A **difference measure** is a function $D : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ which is defined as $D(v_1, v_2) = |v_1 - v_2|$.

Obviously, the difference measure is a distance function since it satisfies non-negativity, symmetry, identity of indiscernibles, and subadditivity. We can use it to define a ranking over probability distributions, and then pick the minimally ranked probability distributions as the aggregation.

Definition 12. Let \preceq be a pre-order relation over \mathcal{P} . A **ranking-based aggregation function**, denoted $F_{\preceq}(T)$, is defined as follows.

$$F_{\preceq}(T) = \{P \in \mathcal{P} \mid \text{for all } P' \in \mathcal{P}, P \preceq P'\}$$

The following definition provides some examples of pre-order relations. For this, \leq_{lex} is the lexicographic ordering over a pair of tuples of numbers N and N' for which there is a permutation of the numbers in rank order from largest to smallest (e.g. $[0.75, 0.5, 0] \leq_{lex} [1, 0.25, 0.25]$).

Definition 13. Let P and P' be probability distributions and let $T = [(\phi_1, v_1), \dots, (\phi_n, v_n)]$ be a profile, and let $[(\phi'_1, v'_1), \dots, (\phi'_n, v'_n)]$ and $[(\phi''_1, v''_1), \dots, (\phi''_n, v''_n)]$ be permutations of T .

$$\begin{aligned} P \preceq_T^{sum} P' &\text{ iff } \sum_{(\phi, v) \in T} D(P(\phi), v) \leq \sum_{(\phi, v) \in T} D(P'(\phi), v) \\ P \preceq_T^{max} P' &\text{ iff } \max_{(\phi, v) \in T} D(P(\phi), v) \leq \max_{(\phi, v) \in T} D(P'(\phi), v) \\ P \preceq_T^{lex} P' &\text{ iff } [D(P(\phi'_1), v'_1), \dots, D(P(\phi'_n), v'_n)] \\ &\leq_{lex} [D(P(\phi''_1), v''_1), \dots, D(P(\phi''_n), v''_n)] \\ P \preceq_T^{more} P' &\text{ iff } |\{i \in \{1, \dots, n\} \mid D(P(\phi_i), v_i) = 0\}| \\ &\geq |\{i \in \{1, \dots, n\} \mid D(P'(\phi_i), v_i) = 0\}| \\ P \preceq_T^{close} P' &\text{ iff } |\{i \in \{1, \dots, n\} \mid D(P(\phi_i), v_i) \leq 0.5\}| \\ &\geq |\{i \in \{1, \dots, n\} \mid D(P'(\phi_i), v_i) \leq 0.5\}| \end{aligned}$$

We explain these rankings as follows. Sum prefers distributions that have the smallest sum difference with the profiles; Max prefers distributions that have the smallest maximum difference with the profiles; Lex prefers distributions that when considering the differences in value for the elements, prefers those that are lexicographically minimal; More prefers distributions that have the most profile elements that are satisfied; and Close is the same as More but allows a margin of error. The Sum, Max, and Lex rankings are inspired by similar rankings for propositional logic by (Konieczny and Pérez 2002). We use these rankings in the ranking-based aggregation function as illustrated next.

Example 8. Consider Table 1, with $\Pi = \{0, 1\}$, $\sigma = \text{pr}$, and profile $T = [(\diamond A, 1), (\diamond B, 1)]$. So $\text{Sat}((\diamond A, 1)) = \{P_1, P_2, P_4\}$ and $\text{Sat}((\diamond B, 1)) = \{P_1, P_3, P_4\}$. For $\preceq \in \{\preceq_T^{max}, \preceq_T^{sum}, \preceq_T^{lex}, \preceq_T^{more}, \preceq_T^{close}\}$, $F_{\preceq}(T) = \{P_1, P_4\}$.

Example 9. Consider Table 1, with $\Pi = \{0, 0.5, 1\}$, $\sigma = \text{gr}$, and profile $T = [(A \wedge \neg B, 1), (A \wedge \neg B, 1), (A \wedge \neg B, 1), (A \wedge \neg B, 0.5), (\neg A \wedge B, 1)]$. For $\preceq \in \{\preceq_T^{max}, \preceq_T^{lex}, \preceq_T^{close}\}$, $F_{\preceq}(T) = \{P_8\}$, and for $\preceq \in \{\preceq_T^{sum}, \preceq_T^{more}\}$, $F_{\preceq}(T) = \{P_2\}$.

The choice of restricted value set can affect the aggregation as illustrated next.

Example 10. Consider Table 1 with $\sigma = \text{gr}$ and the profile $T = [(A \wedge \neg B, 1), (\neg A \wedge B, 1)]$. If $\Pi = \{0, 1\}$, then $F_{\preceq}(T) = \{P_2, P_3\}$, whereas if $\Pi = \{0, 0.5, 1\}$, then $F_{\preceq}(T) = \{P_9, P_0\}$.

Our ranking-based aggregation function satisfies the general properties, and we can partition them using the majority/arbitration properties.

Proposition 8. For $\preceq \in \{\preceq_T^{sum}, \preceq_T^{max}, \preceq_T^{lex}, \preceq_T^{more}, \preceq_T^{close}\}$, the aggregation function $F_{\preceq}(T)$ satisfies viability, consistency, equivalence, fairness, and concordance.

Proposition 9. For $\preceq \in \{\preceq_T^{sum}, \preceq_T^{lex}, \preceq_T^{more}, \preceq_T^{close}\}$, if $\mathcal{G} = \text{Subgraphs}(G^*)$ where G^* is a complete graph, then the aggregation function $F_{\preceq}(T)$ satisfies majority.

Proposition 10. For $\preceq \in \{\preceq_T^{max}\}$, the aggregation function $F_{\preceq}(T)$ satisfies arbitration.

The set of probability distributions returned by an aggregation function provide alternatives for what the probability distribution should be. One way of using this set is to combine the distributions into a single distribution. For this we consider the combination function in the next section.

Combination

A combination function simply takes a set of probability distributions, and returns a probability distribution.

Definition 14. A **combination function** is a function $C : \wp(\mathcal{P}) \rightarrow \mathcal{P}$ such that $C(\{P\}) = P$.

There are various options for the definition of a combination function such as the next definition.

Definition 15. Let $\{P_1, \dots, P_n\} \subseteq \mathcal{P}$ be a set of probability distributions. A **uniform combination** is a combination function C such that for each $G_i \in \mathcal{G}$.

$$P(G_i) = \frac{1}{n} \times \sum_{k \in \{1, \dots, n\}} P_k(G_i)$$

Example 11. Consider Example 10. When $\Pi = \{0, 1\}$, then $F_{\preceq}(T) = \{P_2, P_3\}$. So using the uniform combination, $C(\{P_2, P_3\}) = P_8$

Implicit in the definition of uniform combination is the idea that there is a more refined Π that is prescribed by the input set of probability distributions. For this we introduce the notion of refinement. First, we require a subsidiary function: For a constellation set \mathcal{G} , and a probability distribution $P \in \mathcal{P}$, let $\text{Codomain}(P) = \{v \mid G \in \mathcal{G} \text{ and } P(G) = v\}$.

Definition 16. Let Π be a restricted value set and let $\{P_1, \dots, P_n\} \subseteq \mathcal{P}$ be a set of probability distributions where for each $P_i \in \{P_1, \dots, P_n\}$, $\text{Codomain}(P_i) \subseteq \Pi$. The **refined value set** for Π and $\{P_1, \dots, P_n\}$ is defined as follows where $1/k$ is smallest non-zero value in Π .

$$\text{Refinement}(k, n) = \left\{ \frac{m}{kn} \mid m \in \{1, \dots, kn\} \right\}$$

So for n probability distributions, the refinement of Π will have a cardinality of n times the cardinality of Π . Furthermore, a refinement is a restricted value set.

Proposition 11. If $1/k$ is the smallest non-zero value in restricted value set Π , and $n \in \mathbb{N}$, then $\text{Refinement}(k, n)$ is a restricted value set.

Next we show that if a set of probability distributions take values from a restricted value set, then they can take the same values from its refinement.

Proposition 12. For a set of probability distributions $\{P_1, \dots, P_n\} \subseteq \mathcal{P}$, and restricted value set Π' , such that for all i , $\text{Codomain}(P_i) \subseteq \Pi'$, and $1/k$ is the smallest non-zero value in Π' , if $\text{Refinement}(k, n) = \Pi$, then for all i , $\text{Codomain}(P_i) \subseteq \Pi$.

The following desirable property states that starting with probability distributions that use a finite restricted value set, the resulting distribution uses a finite restricted value set. Whilst this obviously holds for the uniform combination function, it might not hold for interesting alternatives.

Definition 17. A combination function C is **finite** iff for all sets $\{P_1, \dots, P_n\}$, if $C(\{P_1, \dots, P_n\}) = P$, then there is a finite restricted value set Π such that for all i , $\text{Codomain}(P_i) \subseteq \Pi$ and $\text{Codomain}(P) \subseteq \Pi$.

Proposition 13. If C is a uniform combination function, then C is finite.

In the next section, we will consider how we can use aggregation and combination to give a form of induction.

Induction

In general, induction provides a summary of a profile in the form of a probability distribution over a constellations set.

Definition 18. Let \mathcal{T} be a set of profiles and \mathcal{P} be a set of probability distributions. An **induction function** is a function $I : \mathcal{T} \rightarrow \mathcal{P}$ s.t. $I(\emptyset)$ is the uniform distribution, denoted P^{Uniform} , which assigns the same value to every graph in \mathcal{G} .

So an induction function is like aggregation but returns a unique probability distribution. We could take a direct approach by returning the distribution based directly on the

satisfying distributions of each profile element. Essentially, for each subgraph, it is the weighted sum of the belief in that subgraph. However, in this paper, we will take an indirect approach by defining induction in terms of aggregation and combination.

Definition 19. Let $T = [x_1, \dots, x_n]$ be a profile and let $\preceq \in \{\preceq_T^{\text{max}}, \preceq_T^{\text{sum}}, \preceq_T^{\text{more}}, \preceq_T^{\text{lex}}, \preceq_T^{\text{close}}\}$. An **indirect induction function** is an induction function $I : \wp(\mathcal{T}) \rightarrow \mathcal{P}$ such that $I_{\preceq}(T) = P$ where for each $G_i \in \mathcal{G}$

$$P(G_i) = \frac{1}{|F_{\preceq}(T)|} \times \sum_{P_k \in F_{\preceq}(T)} P_k(G_i)$$

Proposition 14. If $I_{\preceq}(T)$ is an indirect induction function, then $I_{\preceq} = C(F_{\preceq}(T))$, where C is the uniform combination function.

Example 12. Consider Table 1. Suppose we have the profile $T = [(A \wedge \neg B, 1), (A \wedge \neg B, 1), (\neg A \wedge B, 1)]$ where $\Pi = \{0, 1\}$, So $\text{Sat}(A \wedge \neg B, 1) = \{P_2\}$, and $\text{Sat}(\neg A \wedge B, 1) = \{P_3\}$. Hence, for $\preceq = \preceq_T^{\text{max}}$, $F_{\preceq}(T) = \{P_2, P_3\}$. Therefore, $I(T) = P$ where $P(G_2) = 1/2$ and $P(G_3) = 1/2$.

Example 13. We now return to the motivating example in Figure 1. Suppose we have the following profile.

$$T = [(\neg A, 1), (A, 1), (A, 1), (A, 0.5)(A \vee B, 1)]$$

For the probability distributions in Table 1, with $\Pi = \{0, 0.5, 1\}$, we get the following by induction.

\preceq	F_{\preceq}	I_{\preceq}
$\preceq_T^{\text{max}}, \preceq_T^{\text{close}}$	$\{P_5, P_7, P_8, P_0\}$	P_{11}
$\preceq_T^{\text{sum}}, \preceq_T^{\text{lex}}$	$\{P_8, P_0\}$	P_{12}
\preceq_T^{more}	$\{P_2, P_3, P_4, P_8, P_0\}$	P_{13}

For this, the refined distributions are defined as follows, with each giving a reasonable summary of the profile.

	G_1	G_2	G_3	G_4
P_{11}	0.25	0.25	0.25	0.25
P_{12}	0	0.25	0.5	0.25
P_{13}	0	0.3	0.4	0.3

For comparison, we give the inferences from each refined distribution to indicate the compromises.

P_{11}	$(\neg A, 0.5)$	$(A, 0.5)$	$(A \vee B, 0.75)$
P_{12}	$(\neg A, 0.5)$	$(A, 0.5)$	$(A \vee B, 1)$
P_{13}	$(\neg A, 0.4)$	$(A, 0.6)$	$(A \vee B, 1)$

Next, we propose some general properties for an induction function (some inspired by (Konieczny and Pérez 2002)). These are the counterparts to those for our aggregation functions, though we do not have counterparts to viability or concordance because induction is a function.

- **itautology:** If $\text{Sat}(T) = \text{Sat}(T)$, then $I(T) = P^{\text{Uniform}}$.
- **iconsistency:** If $\text{Sat}(T) \neq \emptyset$, then $I(T) \in \text{Sat}(T)$.
- **iequivalence:** If $T_1 \equiv T_2$, then $I(T_1) = I(T_2)$.
- **ifairness:** If $\text{Sat}([x_1]) \cap \text{Sat}([x_2]) = \emptyset$, then $I([x_1, x_2]) \notin \text{Sat}(x_2)$.

We explain these properties as follows: itautology captures that when the profile is satisfied by any probability distribution, then it is offering no constraint on the probability distribution and so the uniform distribution will be the

least committing; inconsistency ensures that when the profile is consistent, then the result of induction is one of the satisfying distributions; equivalence ensures that when two profiles are equivalent, then the result of induction is the same; and ifairness ensures when two elements are contradictory, then neither element has one of its satisfying distributions as the result of induction. We consider these properties for ranking-based induction in the next result.

Proposition 15. *For profile T and $\preceq \in \{\preceq_T^{sum}, \preceq_T^{max}, \preceq_T^{more}, \preceq_T^{lex}, \preceq_T^{close}\}$, the induction function $I_{\preceq}(T)$ satisfies itautology, iequivalence, and ifairness. but does not satisfy inconsistency.*

We also introduce two properties that discriminate between two types of induction function. The first aims to minimize the dissatisfaction of the profile as a whole, while the second aims to minimize the dissatisfaction of any individual profile element.

- **imajority:** For any profile T , and probability distribution P , there is a profile element x and $k \in \mathbb{N}$ such that $I(T + [x]^k) = P$.
- **iarbitration:** For any profile T , profile element x , and $k \in \mathbb{N}$, $I(T + [x]^k) = I(T + [x])$.

We can partition our ranking-based induction functions using the majority/arbitration properties.

Proposition 16. *An induction function cannot satisfy both imajority and iarbitration.*

Proposition 17. *For profile T and $\preceq \in \{\preceq_T^{sum}, \preceq_T^{lex}, \preceq_T^{more}, \preceq_T^{close}\}$, if $\mathcal{G} = \text{Subgraphs}(G^*)$ where G^* is a complete graph, the induction function $I_{\preceq}(T)$ satisfies imajority.*

Proposition 18. *For profile T and $\preceq \in \{\preceq_T^{max}\}$, the induction function $I_{\preceq}(T)$ satisfies iarbitration.*

So the difference in the induction functions reflects the difference in the underlying aggregation functions.

Comparison with the Literature

In a sense, our proposal offers a form of learning from data. Each profile element is an item of data, and by generating the probability distribution, we are learning from this data. The probability distribution provides a generalization of the perspectives. For some time, there has been interest in using argumentation for improving machine learning and using machine learning for generating arguments (for a review, see (Cocarascu and Toni 2016)). Out of this literature, there are two recent proposals for learning argument graphs that are based on probabilistic techniques, though as we discuss next, they are different to our proposal.

The first proposal based on probabilistic techniques, by (Riveret and Governatori 2016), is based on labellings for arguments using the usual *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. For the learning, the probability distribution is used to generate labellings that are used as data, and then the argument graph that best describes this data is identified. So probabilistic argumentation is used to

generate the data, but the uncertainty in the argument graph that is learned is not quantified, and so the output is not in the form of a probability distribution over a constellation set.

The second proposal based on probabilistic techniques, by (Kido and Okamoto 2017), takes as input a profile $\langle X_1, \dots, X_n \rangle$ where each X_i is a set of acceptable arguments, and the output is a posterior probability for a set of arguments being the extension. This is calculated using a Bayesian network that incorporates various assumptions about the relationships between choice of semantics and choice of attacks, and how these influence extensions. So the output and underlying mechanisms are completely different to our approach.

There is also another proposal for an argument-based aggregation of reviews (Cocarascu, Rago, and Toni 2019). The reviews are obtained from argumentation dialogues with agents. The aggregation can be used to make predictions using specialized models for aggregating scores for film reviews. However, the aggregation uses the approach of quantitative bipolar framework with gradual semantics, as opposed to a probabilistic approach to argumentation. So, in contrast to our work, the input and output, and underlying mechanisms, do not benefit from a probabilistic approach.

There are some proposals for combining views on the social web, such as (Leite and Martins 2011) which incorporates a form of voting for the structure of the argument graph, and there are a number of proposals in collective argumentation, i.e. how multiple agents combine their argumentation (for a review, see (Bodanza, Tohmé, and Auday 2017)), including aggregating multiple argument graphs (Coste-Marquis et al. 2007; Delobelle et al. 2016), attack aggregation (Tohmé, Bodanza, and Simari 2008), and judgment aggregation for argumentation (Caminada and Pigozzi 2011; Awad et al. 2017), but none of these proposals quantify the uncertainty in the input or output. In contrast, we have uncertainty in the input, and more importantly, we have a probability distribution over the possible argument graphs, instead of a single argument graph as output.

Discussion

In this paper, we have provided a framework for using a profile to generate a probability distribution over a constellation set. Since, there are a number of choices in the framework, we have provided some general properties to guide the choices. Central to our proposal is the need to reason with constellations to determine the probability of dialectical outcomes. Whilst using the constellations approach is computationally expensive (Fazzinga, Flesca, and Parisi 2015), developments in approximation (Fazzinga, Flesca, and Parisi 2013) and automated reasoning could be harnessed (Bistarelli et al. 2018). So in future work we will investigate how we can harness automated reasoning techniques. We will also undertake empirical evaluations including using datasets of online reviews.

References

- Awad, E.; Booth, R.; Tohmé, F.; and Rahwan, I. 2017. Judgement aggregation in multi-agent argumentation. *Jour-*

- nal of Logic and Computation* 27 (1), 227–259 27(1):227–259.
- Bistarelli, S.; Mantadelis, T.; Santini, F.; and Taticchi, C. 2018. Probabilistic argumentation frameworks with metaproblog and conarg. In *Proceedings of ICTAI'18*, 675–679. IEEE Press.
- Bodanza, G.; Tohmé, F.; and Auday, M. 2017. Collective argumentation: A survey of aggregation issues around argumentation frameworks. *Argument and Computation* 8(1):1–34.
- Caminada, M., and Gabbay, D. M. 2009. A logical account of formal argumentation. *Studia Logica* 93:109–145.
- Caminada, M., and Pigozzi, G. 2011. On judgment aggregation in abstract argumentation. *Autonomous Agents and Multi-Agent Systems* 22(1):64–102.
- Cerutti, F.; Norman, T.; Toniolo, A.; and Middleton, S. 2018. Cispaces.org: from fact extraction to report generation computational models of argument: Proceedings of comma 2018 305, 269. In *Proceedings of COMMA'18*, 269–280. IOS Press.
- Cocarascu, O., and Toni, F. 2016. Argumentation for machine learning: A survey. In *Proceedings of COMMA'16*, 219–230. IOS Press.
- Cocarascu, O.; Rago, A.; and Toni, F. 2019. Extracting dialogical explanations for review aggregations with argumentative dialogical agents. In *Proceedings of AAMAS'19*, 1261–1269. IFAAMAS.
- Coste-Marquis, S.; Devred, C.; Konieczny, S.; and Lagasquie-Schiex, M. 2007. On the merging of Dung's argumentation systems. *Artificial Intelligence* 171(10–15):730–753.
- Delobelle, J.; Haret, A.; Konieczny, S.; Mailly, J.; Rossit, J.; and Woltran, S. 2016. Merging of abstract argumentation frameworks. In *Proceedings of KR'16*, 33–42.
- Dung, P., and Thang, P. 2010. Towards (probabilistic) argumentation for jury-based dispute resolution. In *Proceedings of COMMA'10*, volume 216 of FAIA, 171–182. IOS Press.
- Dung, P. M. 1995. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming and n-Person Games. *Artificial Intelligence* 77(2):321–358.
- Fazzinga, B.; Flesca, S.; and Parisi, F. 2013. Efficiently estimating the probability of extensions in abstract argumentation. In *Proceedings of SUM'13*, volume 8078 of LNCS, 106–119.
- Fazzinga, B.; Flesca, S.; and Parisi, F. 2015. On the complexity of probabilistic abstract argumentation frameworks. *ACM Transactions on Computational Logic* 16(3):22.
- Hunter, A., and Thimm, M. 2016. Optimization of dialectical outcomes in dialogical argumentation. *International Journal of Approximate Reasoning* 78:73–102.
- Hunter, A., and Thimm, M. 2017. Probabilistic reasoning with abstract argumentation frameworks. *Journal of Artificial Intelligence Research* 59:565–611.
- Hunter, A.; Polberg, S.; and Thimm, M. 2018. Epistemic graphs for representing and reasoning with positive and negative influences of arguments,. Technical Report arXiv:1802.07489 [cs.AI], Arxiv.
- Hunter, A. 2012. Some foundations for probabilistic abstract argumentation. In *Proceedings of COMMA'12*, volume 245 of FAIA, 117–128. IOS Press.
- Hunter, A. 2013. A probabilistic approach to modelling uncertain logical arguments. *International Journal of Approximate Reasoning* 54(1):47–81.
- Kido, H., and Okamoto, K. 2017. A Bayesian approach to argument-based reasoning for attack estimation. In *Proceedings of IJCAI'17*, 249–255. IJCAI.
- Konieczny, S., and Pérez, R. 2002. Merging information under constraints: a logical framework. *Journal of Logic and Computation* 12(5):773–808.
- Krauthoff, T.; Meter, C.; Baumann, M.; Beta, G.; and Mauve, M. 2018. D-BAS - a dialog-based online argumentation system. In *Proceedings of COMMA'18*, 325–336. IOS Press.
- Leite, J., and Martins, J. 2011. Social abstract argumentation. In *Proceedings of IJCAI'11*, 2287–2292. IJCAI/AAAI.
- Li, H.; Oren, N.; and Norman, T. J. 2011. Probabilistic argumentation frameworks. In *Proceedings of TAFAs'11*, volume 7132 of LNCS, 1–16. Springer.
- Likert, R. 1931. A technique for the measurement of attitudes. *Archives of Psychology* 140:1–55.
- Noor, K.; Hunter, A.; and Mayer, A. 2017. Extraction and analysis of medical arguments from patient experiences expressed on the social web. In *Proceedings of IEA/AIE'2017*, volume 10351 of LNCS, 285–294. Springer.
- Polberg, S., and Hunter, A. 2018. Empirical evaluation of abstract argumentation: Supporting the need for bipolar and probabilistic approaches. *International Journal of Approximate Reasoning* 93:487–543.
- Riveret, R., and Governatori, G. 2016. On learning attacks in probabilistic abstract argumentation. In *Proceedings of AAMAS'16*, 653–661. ACM Press.
- Thimm, M. 2012. A probabilistic semantics for abstract argumentation. In *Proceedings of ECAI'12*, volume 242 of FAIA, 750–755. IOS Press.
- Tohmé, F.; Bodanza, G.; and Simari, G. 2008. Aggregation of attack relations: A social-choice theoretical analysis of defeasibility criteria. In *Proceedings of Foiks'08*, LNCS, 8–23. Springer.