

Probabilistic Argument Graphs for Argumentation Lotteries

Anthony HUNTER^a and Matthias THIMM^b

^a*Department of Computer Science, University College London, UK*

^b*Institute for Web Science and Technology, University of Koblenz-Landau, Germany*

Abstract. Uncertainty about which arguments or attacks should appear in an argument graph means that there is uncertainty as to the structure of the argument graph. When informal arguments are presented, there may be imprecision in the language used, and so the audience may be uncertain as to the structure of the argument graph as intended by the presenter of the arguments. For a presenter of arguments, it is useful to know the audience's argument graph, but the presenter may be uncertain as to the structure of it. To model each of these situations, we can use probabilistic argument graphs. The set of subgraphs of an argument graph is a sample space. A probability value is assigned to each subgraph such that the sum is 1, thereby reflecting the uncertainty over which is the actual subgraph. We can then determine the probability that a particular set of arguments is included or excluded from an extension according to a particular Dung semantics. We harness this to define the notion of an argument lottery, which can be used by the audience to determine the expected utility of a debate, and can be used by the presenter to decide which arguments to present by choosing those that maximize expected utility.

Keywords. Abstract argumentation, Probabilistic argumentation, Utility theory, Lotteries, Audience modelling, Uncertainty in argumentation

1. Introduction

In abstract argumentation, a graph is used to represent a set of arguments and counterarguments. Each node is an argument and each arc from an argument α to an argument β denotes an attack by α on β . It is a well-established and intuitive approach to modelling argumentation, and it offers a valuable starting point for theoretical analysis of argumentation [4]. However, including an argument α in a graph usually means that one is sure that α is a *justifiable* argument, i. e., that it is an argument that makes sense (independently of whether it can be *accepted* after relating it to other arguments). To address the need to represent and reason with (quantified) uncertainty, it has been proposed to use a probability assignment to arguments and to attacks [1]. This can be used to give a probability distribution over the subgraphs of the argument graph, and this can then be used to give a probability assignment for a set of arguments being an admissible set or extension of the argument graph [11,9,10]. The probability distribution over subgraphs denotes the uncertainty over which subgraph is the actual graph that should be used. We refer to an argument graph with a probability distribution over subgraphs as a *probabilistic argument graph*. We believe the following are two important applications for probabilistic argument graphs:

- **From an audience's perspective**, there may be uncertainty as to what the actual argument graph is. 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. There may be uncertainty about which arguments are meant to be attacked by some argument, which occurs frequently when enthymemes are presented. So the audience can collate all the candidates for arguments and attacks, and construct the graph containing them all, and then identify a probability distribution over its subgraphs that reflects their uncertainty about which is the actual graph.
- **From a participant's perspective** (i. e. from the perspective of someone who is about to present arguments and/or attacks to some monological or dialogical argumentation), there may be uncertainty about what the audience regards as the argument graph. 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 an argument graph in mind and the participant will be adding to that graph in the audience's mind. To handle this, the participant may have an argument graph which he/she assumes will subsume the possibilities for the argument graph held by the audience, and then the participant might identify a probability distribution over subgraphs of the argument graph to reflect the uncertainty as judged by the participant over which is the subgraph being used by the audience.

As we will see in this paper, we can investigate probabilistic argument graphs to determine the probability of outcomes of an argument graph. We define these outcomes in the form of a generalization of the notion of extensions that we call *divisions*. A division is a tuple $\langle \Phi, \Psi \rangle$ such that there is an extension of the graph that includes the arguments in Φ and excludes the arguments in Ψ . Using a probabilistic argument graph, we can determine the probability that a tuple is a division. From an audience's perspective, this gives a better understanding of the consequences of the debate that they are observing, and from a participant's perspective, it gives a better understanding of whether s/he will get the desired outcomes from his/her contributions to the argumentation.

We can further exploit probabilistic argument graphs, by introducing the notion of *lotteries* for argumentation. Assume that during a discussion, a debater wants to identify a good argument to bring into the discussion and that the audience of the discussion is considering some subgraph of G as the *true* argument graph. The debater does not know for sure which subgraph is the correct one but he can identify a probability distribution over the subgraphs. Now, suppose he is keen that arguments α and β are accepted by the audience (e. g. they are both in the grounded extension of whichever subgraph the audience is using). So the outcome we want is that α and β are included in the grounded extension. If this is not possible, then perhaps he wants the outcome where α is included and β excluded. Suppose any other outcome is inferior to these two outcomes. By using the probabilistic argument graphs, we are able to determine a probability for each of these outcomes, and we can construct a lottery containing these arguments. If we identify a utility function over outcomes, we can apply utility theory to determine the ex-

pected utility. Furthermore, if we then consider further arguments that we can add to the discussion, we can evaluate the expected utility of each choice of further arguments to put forward. We can then determine which actions (i. e. which arguments to add to the discussion) will offer the maximum expected utility.

The aim of this paper is to develop the use of probabilistic argument graphs, and to apply them to argumentation lotteries. For this, we make the following contributions:

1. Introduce the notion of a division $\langle \Phi, \Psi \rangle$ as an outcome that holds for an argument graph when there is an extension including the arguments in Φ and excluding the arguments in Ψ (Section 3);
2. Introduce the probability that a division $\langle \Phi, \Psi \rangle$ is an outcome for a probabilistic argument graph (Section 4); and
3. Introduce the notion of an argumentation lottery (Section 5) which we show can be used to determine the expected utility of an argument graph, for an audience's perspective, and can be used to enable an agent to determine what would be the best contribution to make in monological or dialogical argumentation in order to maximize its expected utility, for a participant's perspective (Section 6).

We introduce some necessary preliminaries in Section 2 and conclude with a discussion of related work in Section 7. Proofs of technical results can be found in an online appendix¹.

2. Preliminaries

An **abstract argument graph** is a pair $(\mathcal{A}, \mathcal{R})$ where \mathcal{A} is a set and $\mathcal{R} \subseteq \mathcal{A} \times \mathcal{A}$ [4]. 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 $\text{Defended}(\Gamma) = \{\alpha \mid \Gamma \text{ defends } \alpha\}$.

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

In order to present our framework of probabilistic argument graphs we 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 $\text{Sub}(G)$ denote the set of subgraphs of G (i.e. $\{G' \mid G' \sqsubseteq G\}$).

¹http://www.mthimm.de/misc/probarg_comma2014_proofs.pdf




Figure 1. A simple argument graph G and its subgraphs.

Example 1. Consider the argument graph G depicted in Figure 1a and its subgraphs depicted in Figures 1a to 1s. We will use G throughout the paper as a running example.

In the following, we will use the subgraphs of a graph to model uncertainty in the original graph.

3. Divisions of Included and Excluded Arguments

Let $\langle \Phi, \Psi \rangle$ be a tuple of sets of arguments of a graph G , i. e., $\Phi, \Psi \subseteq \text{Nodes}(G)$ and let $\text{Tuples}(G) = \{\langle \Phi, \Psi \rangle \mid \Phi, \Psi \subseteq \text{Nodes}(G)\}$ be the set of all such tuples. A tuple $\langle \Phi, \Psi \rangle \in \text{Tuples}(G)$ is called a *division* of G when there is an extension of G that includes the arguments in Φ and excludes the arguments in Ψ . We use the notion of a division as a generalization of the notion of an extension.

Definition 1. Let G be an argument graph and let $X \in \{\text{co}, \text{pr}, \text{gr}, \text{st}\}$ be a semantics. A tuple $\langle \Phi, \Psi \rangle \in \text{Tuples}(G)$ is a **division** of G w.r.t. X iff there is an extension $E \in \text{Extensions}_X(G)$ such that $\Phi \subseteq E$ and $E \cap \Psi = \emptyset$. Let $\text{Divisions}_X(G)$ be the set of divisions of G w.r.t. X .

Example 2. We continue Example 1. The grounded extension of G is $\{\alpha, \gamma\}$. Hence, $\text{Divisions}_{\text{gr}}(G)$ contains exactly the following divisions of G : $\langle \{\alpha, \gamma\}, \{\beta\} \rangle$, $\langle \{\alpha\}, \{\beta\} \rangle$, $\langle \{\gamma\}, \{\beta\} \rangle$, $\langle \emptyset, \{\beta\} \rangle$, $\langle \{\alpha, \gamma\}, \emptyset \rangle$, $\langle \{\alpha\}, \emptyset \rangle$, $\langle \{\gamma\}, \emptyset \rangle$, $\langle \emptyset, \emptyset \rangle$.

The motivation behind a division is that we might want to know if some arguments are included in an extension and some arguments are excluded in that extension for a given argument graph. For instance, for a graph G containing numerous arguments including arguments α and β , we may want to know whether argument α is included in

and β is excluded from the grounded extension of G . We might be unconcerned about the other arguments in G . Therefore, we want to know if $\langle \{\alpha\}, \{\beta\} \rangle$ is in $\text{Divisions}_{\text{gr}}(G)$.

Proposition 1. For every argument graph G and each semantics $X \in \{\text{co}, \text{pr}, \text{gr}, \text{st}\}$,

1. for every division $\langle \Phi, \Psi \rangle \in \text{Divisions}_X(G)$ it holds that $\Phi \cap \Psi = \emptyset$.
2. for every $\alpha \in \text{Nodes}(G)$, $\langle \{\alpha\}, \{\} \rangle \in \text{Divisions}_X(G)$ or $\langle \emptyset, \{\alpha\} \rangle \in \text{Divisions}_X(G)$.
3. the empty division $\langle \emptyset, \emptyset \rangle$ is always in $\text{Divisions}_X(G)$.
4. for every extension $E \in \text{Extensions}_X(G)$ there is a division $\langle E, \text{Nodes}(G) \setminus E \rangle \in \text{Divisions}_X(G)$.
5. if $\langle \Phi, \Psi \rangle \in \text{Divisions}_X(G)$, $\Phi' \subseteq \Phi$, and $\Psi' \subseteq \Psi$, then $\langle \Phi', \Psi' \rangle \in \text{Divisions}_X(G)$.

Given a graph G and a division $\langle \Phi, \Psi \rangle \in \text{Divisions}_X(G)$ the set of *dividers* of $\langle \Phi, \Psi \rangle$ is the set of subgraphs that have $\langle \Phi, \Psi \rangle$ as a division.

Definition 2. Let G be an argument graph and let $X \in \{\text{co}, \text{pr}, \text{gr}, \text{st}\}$ be a semantics. A graph $G' = (\mathcal{A}', \mathcal{R}') \sqsubseteq G$ is a **divider** for a tuple $\langle \Phi, \Psi \rangle \in \text{Tuples}(G)$ iff $\langle \Phi, \Psi \cap \mathcal{A}' \rangle \in \text{Divisions}_X(G')$. Let $\text{Dividers}_X^G(\langle \Phi, \Psi \rangle)$ be the set of dividers of $\langle \Phi, \Psi \rangle$ w.r.t. G and X .

Example 3. We continue Example 2. There we have among others

$$\text{Dividers}_{\text{gr}}^G(\langle \{\alpha\}, \{\beta\} \rangle) = \{G, G_2, G_3, G_4, G_6, G_9, G_{12}, G_{15}\}$$

$$\text{Dividers}_{\text{gr}}^G(\langle \{\alpha\}, \{\} \rangle) = \{G, G_2, G_3, G_4, G_6, G_7, G_9, G_{11}, G_{12}, G_{15}\}$$

$$\text{Dividers}_{\text{gr}}^G(\langle \{\beta\}, \{\alpha\} \rangle) = \{G_5, G_{10}, G_{14}, G_{16}\}$$

$$\text{Dividers}_{\text{gr}}^G(\langle \{\alpha\}, \{\alpha\} \rangle) = \emptyset$$

$$\text{Dividers}_{\text{gr}}^G(\langle \{\alpha, \beta\}, \{\gamma\} \rangle) = \{G_{11}\}$$

$$\text{Dividers}_{\text{gr}}^G(\langle \{\}, \{\} \rangle) = \{G, G_1, \dots, G_{18}\}$$

Proposition 2. For every argument graph G and each semantics $X \in \{\text{co}, \text{pr}, \text{gr}, \text{st}\}$,

1. If $\Phi \cap \Psi \neq \emptyset$, then $\text{Dividers}_X^G(\langle \Phi, \Psi \rangle) = \emptyset$
2. If $\Phi = \Psi = \emptyset$, then $\text{Dividers}_X^G(\langle \Phi, \Psi \rangle) = \{G' \mid G' \sqsubseteq G\}$
3. If $\Phi \subseteq \Phi'$ and $\Psi \subseteq \Psi'$, then $\text{Dividers}_X^G(\langle \Phi', \Psi' \rangle) \subseteq \text{Dividers}_X^G(\langle \Phi, \Psi \rangle)$

We will use divisions as our outcomes when we consider lotteries. For this, we need to ensure that each outcome is disjoint from the other outcomes. For this, we use the following definition.

Definition 3. Let G be an argument graph, and let $T, T' \in \text{Tuples}(G)$. T and T' are **disjoint** for G and w.r.t. semantics X iff $\text{Dividers}_X^G(T) \cap \text{Dividers}_X^G(T') = \emptyset$. A set $\mathcal{T} \subseteq \text{Tuples}(G)$ is **pairwise disjoint** iff for each $T, T' \in \mathcal{T}$, T and T' are disjoint.

Example 4. We continue Example 3 and consider grounded semantics. Then $\langle \{\alpha\}, \{\beta\} \rangle$ and $\langle \{\beta\}, \{\alpha\} \rangle$ are disjoint and $\langle \{\alpha\}, \emptyset \rangle$ and $\langle \{\gamma\}, \emptyset \rangle$ are not disjoint as, e. g., G_{12} is a divider for both of them.

Proposition 3. Tuples $T = \langle \Phi_1, \Psi_1 \rangle, T' = \langle \Phi_2, \Psi_2 \rangle \in \text{Tuples}(G)$ with $\Phi_1 \cap \Psi_1 = \emptyset$ and $\Phi_2 \cap \Psi_2 = \emptyset$ are disjoint iff either $\Psi_1 \cap \Phi_2 \neq \emptyset$ or $\Psi_2 \cap \Phi_1 \neq \emptyset$.

Definition 4. A set of tuples $\{T_1, \dots, T_k\} \subseteq \text{Tuples}(G)$ is **exhaustive** for G w.r.t. semantics X iff $\text{Dividers}_X^G(T_1) \cup \dots \cup \text{Dividers}_X^G(T_k) = \{G' \mid G' \sqsubseteq G\}$.

Example 5. We continue Example 4. The following set of tuples $\{T_1, \dots, T_8\}$ is exhaustive for G with respect to grounded semantics:

$$\begin{aligned} T_1 &= \langle \{\alpha, \beta, \gamma\}, \{\} \rangle & T_2 &= \langle \{\alpha, \beta\}, \{\gamma\} \rangle & T_3 &= \langle \{\alpha, \gamma\}, \{\beta\} \rangle & T_4 &= \langle \{\beta, \gamma\}, \{\alpha\} \rangle \\ T_5 &= \langle \{\alpha\}, \{\beta, \gamma\} \rangle & T_6 &= \langle \{\beta\}, \{\alpha, \gamma\} \rangle & T_7 &= \langle \{\gamma\}, \{\alpha, \beta\} \rangle & T_8 &= \langle \{\}, \{\alpha, \beta, \gamma\} \rangle \end{aligned}$$

Furthermore, the tuples are pairwise disjoint (i. e. for each $i, j \in \{1, \dots, 8\}$, T_i and T_j are disjoint for grounded semantics).

Proposition 4. Let G be an argument graph and define a set of divisions D via

$$D = \{ \langle \Phi, \Psi \rangle \mid \Phi \subseteq \text{Nodes}(G) \text{ and } \Psi = \text{Nodes}(G) \setminus \Phi \}$$

Then D is exhaustive for G w.r.t. any semantics X and D is pairwise disjoint for G w.r.t. grounded semantics.

Definition 5. A division T **subsumes** a set of divisions $\{T_1, \dots, T_i\}$ iff

$$\text{Dividers}_X^G(T_1) \cup \dots \cup \text{Dividers}_X^G(T_i) = \text{Dividers}_X^G(T)$$

Proposition 5. Let D be an exhaustive set of divisions where the divisions in D are pairwise disjoint. Assume T is a division not in D . Also let $D' \subseteq D$ be a subset of divisions. If T subsumes D' , then the set of divisions $(D \setminus D') \cup \{T\}$ is exhaustive and pairwise disjoint.

Example 6. We continue Example 5. Recall that $\{T_1, \dots, T_8\}$ is exhaustive and pairwise disjoint for G with respect to grounded semantics. The tuple $T_9 = \langle \{\beta\}, \{\} \rangle$ subsumes $\{T_1, T_2, T_4, T_6\}$. Therefore $\{T_3, T_5, T_7, T_8, T_9\}$ is exhaustive and pairwise disjoint for G with respect to grounded semantics.

Divisions provide a useful generalization of the notion of extensions which we will harness when we consider probability distributions over subgraphs and lotteries in the following subsections.

4. Probability distributions

Given an argument graph G we represent the uncertainty we may have over the arguments and/or attacks by using the set of subgraphs $\text{Sub}(G)$ as the sample space. This means we are unsure which subgraph is the “correct” subgraph. Then using this sample space, we define a probability distribution as follows. Note, in previous work, we restricted consideration to the *spanning* subgraphs—i. e. subgraphs containing all attacks on a subset of arguments—thereby denoting uncertainty in the arguments [9], or on the *full* subgraphs—i. e. subgraphs containing all arguments but a subset of attacks—thereby denoting uncertainty in the attacks [10]. Here we allow the representation of uncertainty in both arguments and attacks.

Definition 6. Let G be an argument graph. A **probability distribution** P for G is a function $P : \{G' \mid G' \sqsubseteq G\} \rightarrow [0, 1]$ such that $\sum_{G' \sqsubseteq G} P(G') = 1$.

Given a probability distribution over subgraphs, we can obtain the probability of each argument and each attack as a marginal distribution. More specifically, for an argument graph G and a probability distribution P , the **marginal distribution for an argument** α is $P(\alpha) = \sum_{G' \sqsubseteq GS.t.\alpha \in \text{Nodes}(G')} P(G')$. The **marginal distribution for an attack** (α, β) is $P((\alpha, \beta)) = \sum_{G' \sqsubseteq GS.t.(\alpha, \beta) \in \text{Arcs}(G')} P(G')$. Note that these probabilities describe the uncertainty to which an argument or attack is believed to be justifiable, i. e. whether it is appropriate to consider this element to be present in the argument graph. In particular, a high probability of an argument does not necessarily imply that the argument is highly acceptable, see below.

Example 7. We continue Example 6. Define a probability distribution P on G via $P(G) = P(G_5) = P(G_8) = 0.1$, $P(G_9) = 0.7$, and $P(G') = 0$ for the remaining subgraphs G' of G . Then the marginal distributions are as follows: $P(\alpha) = 1$, $P(\beta) = 1$, $P(\gamma) = 0.2$, $P((\alpha, \beta)) = 0.9$, $P((\beta, \alpha)) = 0.3$, and $P((\gamma, \beta)) = 0.1$.

Now we use the probability distribution over subgraphs to give a probability that a division holds. As defined below, the probability that a tuple $T \in \text{Tuples}(G)$ is a division of G is the sum of the probabilities of the subgraphs for which T is a division. So the probability that T is a division of G is zero when there is no subgraph G for which T is a division.

Definition 7. Let G be an argument graph and let $X \in \{\text{co}, \text{pr}, \text{gr}, \text{st}\}$ be a semantics. Also let P be a probability distribution of G . For a tuple $\langle \Phi, \Psi \rangle \in \text{Tuples}(G)$, the **probability of a division** w.r.t. X is

$$P_X(\langle \Phi, \Psi \rangle) = \sum_{G' \in \text{Dividers}_X^G(\langle \Phi, \Psi \rangle)} P(G')$$

Example 8. We continue Example 7. There we have $P(\langle \{\alpha\}, \{\beta\} \rangle) = 0.8$ (as G and G_9 are the only dividers with positive probabilities) and $P(\langle \{\beta\}, \{\alpha\} \rangle) = 0.1$ (as G_5 is the only divider with positive probability).

The notion of the probability of a division subsumes the definition of the probability that a set of arguments is an extension, and it subsumes the definition for the probability that an argument is an inference, cf. [11,9,10].

Proposition 6. For every argument graph G and each semantics $X \in \{\text{co}, \text{pr}, \text{gr}, \text{st}\}$,

1. If $\Phi \cap \Psi \neq \emptyset$, then $P_X(\langle \Phi, \Psi \rangle) = 0$
2. If $\Phi = \Psi = \emptyset$, then $P_X(\langle \Phi, \Psi \rangle) = 1$
3. If $\Phi \subseteq \Phi'$ and $\Psi \subseteq \Psi'$, then $P_X(\langle \Phi', \Psi' \rangle) \leq P_X(\langle \Phi, \Psi \rangle)$

5. Argumentation as a lottery

We start by briefly reviewing the notion of a lottery. A **lottery** is a probability distribution over a set of possible outcomes. A lottery with possible outcomes ϕ_1, \dots, ϕ_n that can

occur with probabilities p_1, \dots, p_n is written $[p_1, \phi_1; \dots; p_n, \phi_n]$. For a utility function U , the **expected utility of a lottery** L , denoted $E(L, U)$, is given by $E(L, U) = \sum_{i=1}^n p_i U(\phi_i)$. We can view an argument graph G as invoking a lottery. For that we use divisions as outcomes and the probability of a division holding as the probability of the outcome. Furthermore, it is quite natural to think of divisions (i.e. inclusions and exclusions of arguments) as having utility. For example, for an argument graph containing arguments α , β , and γ , and we prefer to have α and β and to not have γ , otherwise we prefer either α or β and not γ , otherwise we are indifferent about the outcome, then we have the preferences over outcomes where $\langle \{\alpha, \beta\}, \{\gamma\} \rangle$ is most preferred, $\langle \{\alpha\}, \{\beta, \gamma\} \rangle$ and $\langle \{\beta\}, \{\alpha, \gamma\} \rangle$ are the second most preferred, and then $\langle \{\}, \{\alpha, \beta\} \rangle$ is the least preferred. Since, we can identify this preference ordering, we can identify a utility function to indicate the degree to which we prefer each of the options. For instance, we could let the utility function U be $U(\langle \{\alpha, \beta\}, \{\gamma\} \rangle) = 10$, $U(\langle \{\alpha\}, \{\beta, \gamma\} \rangle) = 5$, $U(\langle \{\beta\}, \{\alpha, \gamma\} \rangle) = 5$ and $U(\langle \{\}, \{\alpha, \beta\} \rangle) = 0$.

We formalize the construction of a lottery for argumentation as follow. For this we need to ensure that the dividers we use as outcomes in the lottery as pairwise disjoint and together they are exhaustive.

Definition 8. . Let G be the argument graph, let $\{T_1, \dots, T_k\}$ be a set of divisions, and let P be a probability distribution. The tuple $[P_X(T_1), T_1; \dots; P_X(T_k), T_k]$ is an **argumentation lottery** for G w.r.t. semantics X iff

1. $\sum_{T_i \in \{T_1, \dots, T_k\}} P_X(T_i) = 1$
2. $\{T_1, \dots, T_k\}$ is exhaustive and pairwise disjoint for G w.r.t. X

Example 9. We continue Example 8. Recall that $\{T_3, T_5, T_7, T_8, T_9\}$ is exhaustive and pairwise disjoint for G with respect to grounded semantics (see Example 6) and let P be as defined in Example 7. This gives the following argumentation lottery

$$[P(T_3), T_3; P(T_5), T_5; P(T_7), T_7; P(T_8), T_8; P(T_9), T_9] = [0.1, T_3; 0.7, T_5; 0, T_7; 0.1, T_8; 0.1, T_9]$$

Define a utility function U via $U(T_3) = 10$, $U(T_5) = 5$, $U(T_7) = 5$, $U(T_8) = 0$, and for $k \in \{1, 2, 4, 6, 9\}$, $U(T_k) = -10$. Observe that U favours α and/or γ , but not β , in our grounded extension. Therefore the expected utility is $(0.1 \cdot 10) + (0.7 \cdot 5) + (0.0 \cdot 5) + (0.1 \cdot 0) + (0.1 \cdot -10) = 3.5$.

The following result shows that for grounded semantics there is always an argumentation lottery. Furthermore, with the use of subsumption, we can restructure the argumentation lottery to reduce the number of outcomes as illustrated in the above example.

Proposition 7. Let G be an argument graph, P a probability distribution on G , and $\{T_1, \dots, T_k\} = \{ \langle \Phi, \Psi \rangle \mid \Phi \subseteq \text{Nodes}(G), \Psi = \text{Nodes}(G) \setminus \Phi \}$. Then $[P(T_1), T_1; \dots; P(T_k), T_k]$ is an argumentation lottery for G with respect to grounded semantics.

For other semantics such as preferred and stable we can also construct argumentation lotteries. However, we are not able to just resort to $D = \{ \langle \Phi, \Psi \rangle \mid \Phi \subseteq \text{Nodes}(G) \text{ and } \Psi = \text{Nodes}(G) \setminus \Phi \}$. Rather, we need to consider the actual argument graph to determine which divisions are disjoint.

We believe that expected utility is a useful formal tool for an audience to judge argumentation. From the audience's perspective, we are interested in modelling how a

member of the audience may evaluate some arguments. For example, a member of the audience of a political speech may listen to the arguments and counterarguments that the politician has presented, or a member of the audience of a debate may hear the arguments and counterarguments exchanged by the participants. In each case, an argument graph is produced. The member of the audience then may look at the arguments and the attacks and she may be uncertain whether some of the arguments should be included in the graph (perhaps some arguments are rephrasing of previously expressed arguments), and/or whether some of the attacks hold (perhaps the arguments are enthymemes, and she doubts that the enthymemes can be decoded so that it can be attacked by the given counterarguments). In order to represent the uncertainty in the arguments and attacks, the member of the audience identifies a probability distribution over the subgraphs. With this probability distribution, she can determine the probability that specific arguments are included or excluded according to specific semantics. Furthermore, by determining the expected utility of the argumentation lottery, she can determine the worth of the consequences of the debate to her in utility-theoretic terms.

6. Maximizing expected utility in argumentation

We now consider the argumentation lottery from the perspective of the participant. When an agent presents an argument α , this can be viewed as a lottery by the agent since there is uncertainty about whether α will be included or excluded from the viewpoint of the audience according to some semantics. If the agent's probability distribution is P then it can assess the outcome of presenting argument α by evaluating the lottery

$$[P(\langle\{\alpha\}, \emptyset\rangle), \langle\{\alpha\}, \emptyset\rangle; P(\langle\emptyset, \{\alpha\}\rangle), \langle\emptyset, \{\alpha\}\rangle]$$

with respect to its utility function U . Now suppose the agent has a choice of arguments to present say α_1 , α_2 , or α_3 , and for each of the arguments $\alpha_i \in \{\alpha_1, \alpha_2, \alpha_3\}$, if α_i is presented, the agent is unsure whether α_i will be in, out or undecided from the viewpoint of the audience according to, e. g., grounded semantics. So each α_i is an option for an action with an associated lottery L_i , respectively, of the above form. Given these lotteries, we can choose the argument α_i that maximizes expected utility. In the same way, suppose that an agent has a choice of which *sets of arguments* to present, say A_1 , A_2 , or A_3 , but the agent is concerned about another argument β in G . For instance, in a dialogue, β may have been given earlier, and the agent wants to know which would be the best arguments to add at this stage in order to get a particular outcome concerning β .

We organize these ideas as follows. Let the **proponent** be the person who wants to make a contribution to a discussion (or debate, etc). A **contribution** is one or more arguments and attacks. Suppose the proponent can choose between a number of options for the contributions. So each **option** for a contribution contains one or more arguments. Let the options be C_1, \dots, C_k . Now suppose G is the argument graph that includes all the possible arguments and attacks that the intended audience may currently entertain, called the **current argument graph**. For each option for a contribution $C_i \in \{C_1, \dots, C_k\}$, let $G + C_i$ be the argument graph obtained by augmenting G with the arguments and attacks in C_i . Suppose for each argument graph $G + C_i$, the proponent has a probability function P^i over the subgraphs of $G + C_i$. Furthermore, suppose for each argument graph

$G + C_i$, the proponent has some arguments it wants included or excluded. In other words, the proponent can identify an exhaustive and pairwise disjoint set of divisions for each argument graph $G + C_i$, together with an associated utility function. Therefore, for each argument graph $G + C_i$, the proponent can produce an argumentation lottery L_i for the action of making the contribution C_i .

So there are options for contributions C_1, \dots, C_k , and for each option for a contribution $C_i \in \{C_1, \dots, C_k\}$, there is a lottery L_i . Hence, the expected utility of each lottery can be calculated as described in general above. According to utility theory, the best option for a contribution is the contribution C_j for which the expected utility of the lottery L_j is maximal. We formalize this process in the following definition of a game. Note, for this, the lottery is only implicit in the definition.

Definition 9. A game is tuple (G, P, U, D, O) where $G = (\mathcal{A}, \mathcal{R})$ is an argument graph, P is a probability distribution over subgraphs of G , U is a utility function over divisions, D is a set of divisions that is exhaustive and pairwise disjoint, and O is a set of contributions. For each $C^i \in O$,

- C^i is a contribution $(\mathcal{B}^i, \mathcal{S}^i)$ with arguments \mathcal{B}^i and attacks $\mathcal{S}^i \subseteq (\mathcal{A} \cup \mathcal{B}^i \times \mathcal{A} \cup \mathcal{B}^i)$.
- P^i is a probability distribution over subgraphs of $G + C^i$ such that for each $G' \sqsubseteq G$, $P^i(G' + C^i) = P(G')$.

For each $C^i \in O$ and $X \in \{\text{co}, \text{pr}, \text{gr}, \text{st}\}$ define $E(G + C^i) = \sum_{T_j \in D} P_X^i(T_j) \times U(T_j)$. A contribution C^i which maximizes $E(G + C^i)$ is an **optimal contribution** to G .

In the above definition, we do not explicitly construct a lottery but we calculate the expected utility directly from the outcomes (i. e. the divisions specified in O) using the probability distribution P^i and the utility function U .

Example 10. Consider the argument graph G_1 and its subgraphs with non-zero probability given in Figure 2a, together with associated probability distribution. Assume that the outcomes are $\{\{\alpha\}, \{\}\}$ and $\{\{\}, \{\alpha\}\}$ where $U(\{\{\alpha\}, \{\}\}) = 10$ and $U(\{\{\}, \{\alpha\}\}) = -10$. Hence, the expected utility of the argumentation lottery is $(0.2 \times 10) + (0.8 \times -10) = -6$.

Now consider the contribution $C^1 = (\{\delta\}, \{(\delta, \beta)\})$ which we use to give the graph $G_1^1 = G_1 + C^1$ together with the subgraphs with non-zero probability depicted in Figure 2b. Hence, the expected utility of the argumentation lottery is $(0.8 \times 10) + (0.2 \times -10) = 6$.

Finally consider the contribution $C^2 = (\{\varepsilon\}, \{(\varepsilon, \gamma)\})$ which we use to give the graph $G_1^2 = G_1 + C^2$ together with the subgraphs with non-zero probability depicted in Figure 2c. Hence, the expected utility of the argumentation lottery is $(0.3 \times 10) + (0.7 \times -10) = 4$.

So both contributions turn a negative expected utility into a positive expected utility, with C^1 being the contribution that maximizes utility.

So by harnessing a probabilistic argumentation graph, and an argumentation lottery, a participant can optimize its choice of actions in argumentation. This can be for example in monological argumentation (e. g. in a speech or a written article) or dialogical argumentation (e. g. a discussion or debate) when a participant wants to present argu-

G_1	$\beta \rightarrow \alpha \leftarrow \gamma$	0.1	G_1^1	$\delta \rightarrow \beta \rightarrow \alpha \leftarrow \gamma$	0.1	G_1^2	$\beta \rightarrow \alpha \leftarrow \gamma \leftarrow \varepsilon$	0.1
G_2	$\beta \rightarrow \alpha$	0.6	G_2^1	$\delta \rightarrow \beta \rightarrow \alpha$	0.6	G_2^2	$\beta \rightarrow \alpha \quad \varepsilon$	0.6
G_3	$\alpha \leftarrow \gamma$	0.1	G_3^1	$\delta \quad \alpha \leftarrow \gamma$	0.1	G_3^2	$\alpha \leftarrow \gamma \leftarrow \varepsilon$	0.1
G_4	α	0.2	G_4^1	$\delta \quad \alpha$	0.2	G_4^2	$\alpha \quad \varepsilon$	0.2

(a)

(b)

(c)

Figure 2. Subgraphs and probability distributions for Example 10

ments and/or counterarguments in order to convince the audience that some particular arguments should be accepted and some should be rejected.

We implemented the concepts introduced in this paper and, in particular, the game setting described in this section in the Tweety library for Artificial Intelligence [16]. Moreover, we conducted some empirical evaluation of the lottery approach for move selection in order to validate its feasibility and performance compared to a simple utility-based move selection². More precisely, we generated random argument graphs (with an average of 5 arguments and 0.3 attack probability between any two distinct arguments, no self-attacks), a random probability distribution on the subgraphs, and a random utility function (with all utilities in the range $[0, 1]$). Then we sampled (wrt. to the probability distribution) some subgraph to be the actual subgraph assigned to the audience. Afterwards we used our lottery-based approach to determine the set of arguments that should be brought forward and computed the resulting utility after the audience incorporated this set into its argument graph. As a baseline approach, we also determined another set of arguments by only maximizing utility and without taking the probability function and lotteries into account—i. e. the contribution which resulted in the graph with maximum utility was chosen—, and computed the resulting utility on the actual graph. We repeated this experiment 500 times and obtained an average utility for the baseline approach of approximately 0.5571 (standard deviation 0.0953) and an average utility for the lottery approach of approximately 0.6494 (standard deviation 0.0657; the difference is significant with 99% confidence). This (preliminary) investigation shows already the feasibility of our approach. For future work, we aim at conducting further and more systematic experiments.

7. Discussion

In this paper, we have investigated how probabilistic argumentation can be harnessed to formalize the notion of a lottery for argumentation. An argumentation lottery can be used to judge the expected utility of the outcomes in an argument graph. Furthermore, an argumentation lottery can be constructed for each of a number of possible contributions that can be made to a discussions, debate, etc. These lotteries can then be used to determine the contribution that maximizes the expected utility. Therefore, an agent making a decision on what contribution (if any) to make in argumentation now has a formal tool to make the best choice. We can consider uncertainty from either the audience’s perspective or the participant’s perspective. In previous work, we have modelled audiences in terms

²<http://tiny.cc/TweetyLottery>

of the beliefs and desires to assist a participant in choosing the most believable or the most desirable arguments to make [8,7]. However, that did not consider the uncertainty associated with modelling the audience, and it did not provide a utility-theoretic framework. Utility theory has been considered in other frameworks for argumentation (for example [12,14,13]). Moreover, integrating quantitative uncertainty into argumentation theory has also been investigated in, e. g., [11,5,15]. However, this is the first paper that provides a framework for using probabilistic argumentation for argumentation lotteries.

Our work has strong relationships to strategic issues in multi-agent argumentation [14,13,6] where also game theoretical means are utilized for deciding on the next best move in dialogues. In future work, we would like to investigate how our approach could be used in such settings. We would also like to investigate how our approach could be integrated with techniques for updating argument graphs to enforce particular outcomes (see for example [3,2]) since our notion of a contribution can be regarded as an update.

References

- [1] H. Barringer, D. Gabbay, and J. Woods. Temporal Dynamics of Support and Attack Networks: From Argumentation to Zoology. In *Mechanizing Mathematical Reasoning*, pages 59–98. Springer, 2005.
- [2] R. Baumann. What Does it Take to Enforce an Argument? Minimal Change in abstract Argumentation. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI'12)*, pages 127–132. IOS Press, 2012.
- [3] C. Cayrol, D. F. and M. Lagasquie-Schiex. Change in Abstract Argumentation Frameworks: Adding an Argument. *Journal of Artificial Intelligence Research*, 38:49–84, 2010.
- [4] P. Dung. On the Acceptability of Arguments and its Fundamental Role in Nonmonotonic Reasoning, Logic Programming, and N-Person Games. *Artificial Intelligence*, 77:321–357, 1995.
- [5] P. Dung and P. Thang. Towards (Probabilistic) Argumentation for Jury-based Dispute Resolution. In *Proceedings of the Third International Conference on Computational Models of Argument (COMMA'10)*, pages 171–182. IOS Press, 2010.
- [6] C. Hadjilokis, Y. Siantos, S. Modgil, E. Black, and P. McBurney. Opponent Modelling in Persuasion Dialogues. In F. Rossi, editor, *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI'13)*, pages 164–170, 2013.
- [7] A. Hunter. Making Argumentation more Believable. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI'04)*, pages 269–274. MIT Press, 2004.
- [8] A. Hunter. Towards Higher Impact Argumentation. In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI'04)*, pages 275–280. MIT Press, 2004.
- [9] A. Hunter. Some Foundations for Probabilistic Abstract Argumentation. In *Proceedings of the Fourth International Conference on Computational Models of Argument (COMMA'12)*, pages 117–128, 2012.
- [10] A. Hunter. Probabilistic Qualification of Attack in Abstract Argumentation. *International Journal of Approximate Reasoning*, 55:607–638, 2013.
- [11] H. Li, N. Oren, and T. Norman. Probabilistic Argumentation Frameworks. In *Proceedings of the First International Workshop on the Theory and Applications of Formal Argumentation (TAFa'11)*, 2011.
- [12] I. Rahwan and K. Larson. Argumentation and Game Theory. In I. Rahwan and G. R. Simari, editors, *Argumentation in Artificial Intelligence*, pages 321–339. Springer, 2009.
- [13] T. Rienstra, M. Thimm, and N. Oren. Opponent Models with Uncertainty for Strategic Argumentation. In F. Rossi, editor, *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI'13)*, pages 332–338, 2013.
- [14] R. Riveret, H. Prakken, A. Rotolo, and G. Sartor. Heuristics in argumentation: A game theory investigation. In *Computational Models of Argument (COMMA 2008)*, pages 324–335. IOS Press, 2008.
- [15] M. Thimm. A probabilistic semantics for abstract argumentation. In *Proceedings of the European Conference on Artificial Intelligence (ECAI'12)*, pages 750–755, 2012.
- [16] M. Thimm. Tweety - A Comprehensive Collection of Java Libraries for Logical Aspects of Artificial Intelligence and Knowledge Representation. In *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning (KR'14)*, 2014.