A Bayesian Probabilistic Argumentation Framework for Learning from Online Reviews

Kawsar Noor
Department of Computer Science
University College London
London, United Kingdom
kawsar.noor.15@ucl.ac.uk

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

***Abstract*—In the real world it is common for agents to posit arguments concerning an issue but not directly specify the attack relations between them. Nonetheless the agent may have these attacks in mind and instead they may provide a proxy indicator through which one can infer the agent's intended argument graph (arguments and attacks). Consider online reviews, where reviews are collections of arguments for and against the product (positive and negative) under review and the rating indicates whether the positive or negative arguments succeed ultimately. In previous work [1] we have proposed a method that formalises this intuition and uses the constellations approach to probabilistic argumentation to construct a probability distribution over the set of arguments graphs the agent may have had in mind. In this paper we extend this proposal and provide a method, that uses Bayesian inference, to update the initial probability distribution using real data. We evaluate our proposal by conducting a number of simulations using synthetic data.**

***Index Terms*—abstract argumentation, probabilistic argumentation**

I. Introduction

A situation involving argumentation can be modelled at a high-level by a directed graph. Each node denotes an argument, and each arc denotes an attack by one argument on another. Proposals have made for principled ways to identify subsets of the arguments in the graph that can be regarded as winning arguments [2]. These subsets are called extensions and they constitute coalitions of arguments that are acceptable together.

In this paper we use argumentation in a setting in which agents propose a set of arguments but do not provide the attacks between them. Instead they provide additional data that can be used to infer these attacks and thus identify the agent's intended graph. To illustrate consider online product reviews where each review can be thought of as a collection of arguments for (positive arguments) and against the product (negative arguments). These reviews, in many cases, are accompanied by ratings (numeric or graded). We interpret these ratings as a proxy though which we can understand the probability that a set of arguments (positive or negative) win; thus if the rating is highly positive then we assume the positive arguments are likely accepted and vice versa.

In our previous work [1] we have proposed a method for identifying a probability distribution over the set of argument graphs that the agent may have had in mind. The approach uses the constellations approach to probabilistic argumentation [10] [11] [6] which allows one to model uncertainty over the topology over the graph; more precisely uncertainty over a set (constellation) of argument graphs. This initial distribution captures our belief that an agent intended a particular graph based solely on the proxy indicator they have provided. It represents the best informed inference we can make about an agent's intended graph in the absence of the the actual graph or a human annotator to attempt to identify the graph for us. In Bayesian terms this represents our prior distribution.

When data about what each agent intended is acquired, e.g. through annotation, we can then update our initial distribution using the data. For this we propose using Bayesian inference which allows us to generate new beliefs (posterior distribution) by incorporating observed/ real data into our prior belief.

This paper is structured as follows. In section II we introduce argumentation in the context of reviews and give an overview of how Bayesian inference can be applied in this context. In sections III and IV we define the prior and likelihood functions for Bayesian inference and finally in section V we demonstrate our proposal by conducting a number of simulations using synthetic data.

II. Modelling Views

An argument graph is a tuple $G = (A, R)$ where A denotes a set of arguments and $R \subseteq A \times A$ denotes a set of binary attack relations between the arguments. Extension-based semantics are semantics that can be used to identify sets of collectively coherent arguments in an argument graph. At the heart of these semantics is the property of conflict-freeness which states that a set of arguments $B \subseteq A$ is **conflict-free** iff no two arguments $a, b \in B$ exists s.t $(a.b) \in R$. An argument b is **defended** by a set $B \in A$ iff any argument $a \in A$ attacks b then $\exists c \in B$ s.t. $(c, a) \in R$. A conflict-free set $B \subseteq A$ is an **admissible** extension iff each argument in B is defended by B. An admissible extension B is an **complete** extension iff each argument defended by B is in B. A complete extension B is a **grounded** extension if it is minimal (w.r.t set inclusion).

In the context of this paper we use argument graphs to model conflict between arguments in reviews. We define a review as having two components; a textual component and

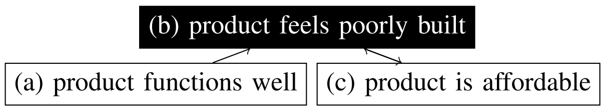

Fig. 1: An example of an argument graph containing two positive arguments (a,c) and one negative argument (b)

a numerical rating. The textual component is comprised of arguments for the product (positive) and against it (negative).

Definition 1. *Let A^+ be a set of positive arguments and A^- be a set of negative arguments s.t. $A^+ \cap A^- = \emptyset$. Let the minimum rating be b^{Neg}_{min} and the maximum be b^{Pos}_{max}. A* ***view*** *is a tuple $v = (A, b)$ where $A \subseteq A^+ \cup A^-$ and $b \in [b^{\mathsf{Neg}}_{min}, b^{\mathsf{Pos}}_{max}]$ is a* ***rating*** *s.t $b^{\mathsf{Neg}}_{min}, b^{\mathsf{Pos}}_{max} \in \mathbb{R}$ and $b^{\mathsf{Neg}}_{min} < b^{\mathsf{Pos}}_{max}$.*

Example 1. *Consider the arguments depicted in Figure 1 and rating . Some examples of views using the arguments $\{a, b, c\}$ with a rating scale $R = \{1..10\}$ would be $(\{a, b, c\}, 9)$ and $(\{a, b\}, 10)$*

When considering the set of possible argument graphs that can be attributed to a view we consider all the bipartite graphs that can be constructed using the positive and negative arguments in the view. We refer to this set as the **graph space**. Formally we say that given disjoint sets A^+, A^- that the graph space is the set returned by the function $\mathsf{Space}(A^+, A^-) = \{(A^+ \cup A^-, R) | R \in \mathcal{P}((A^+ \times A^-) \cup (A^- \times A^+)\}$. An example of a graph space given 2 positive and 1 negative argument is provided in Table I.

In order to identify a probability distribution over the graph space for a view we have previously proposed a method [1] for this; underlying our approach is the assumption that the rating indicates which set of arguments (positive/negative) are acceptable/winning. In Bayesian terms this distribution represents our prior as it captures our beliefs about the review before accounting for any real data.

When real graphs concerning the reviews are acquired, e.g. via annotation, then this distribution needs to be updated to reflect the data. Such annotations we refer to as observations. Formally we say that an observation is the assignment of an argument graph to a view.

Definition 2. *Let $v = (A, b)$ be a view where $A \subseteq A^+ \cup A^-$. An* ***observation*** *is a tuple $t = (G, b)$ where $G \in \mathsf{Space}(A^+, A^-)$.*

Example 2. *Continuing with the example in Figure 1 an example observation would then be $((\{a, b, c\}, \{(a, b)\}), 9)$*

Given a set of observations we then have the necessary components for Bayesian inference.

Definition 3. *Given a set of positive and negative arguments A^+ and A^-, a view $v = (A, b)$ s.t $A \subseteq A^+ \cup A^-$, a graph space $S = \mathsf{Space}(A^+, A^-)$ and a set of of observations T we say that our updated belief is given by the function:*

No	Graph	Agg	P(G) for rating		
			10	7	4
G1	a → b ← c	2	0.16	0	0
G2	a b ← c	1.25	0.13	0.01	0.01
G3	a → b c	1.25	0.13	0.01	0.01
G4	a → b ↔ c	1.25	0.13	0.01	0.01
G5	a ↔ b ← c	1.25	0.13	0.01	0.01
G6	a → b → c	0.5	0.08	0.05	0.02
G7	a ← b ← c	0.5	0.08	0.05	0.02
G8	a b ↔ c	-0.25	0.05	0.11	0.05
G9	a ↔ b c	-0.25	0.05	0.11	0.05
G10	a b c	-0.25	0.05	0.11	0.05
G11	a ↔ b ↔ c	-1	0.02	0.15	0.08
G12	a b → c	-1.75	0.01	0.08	0.13
G13	a ← b c	-1.75	0.01	0.08	0.13
G14	a ↔ b → c	-1.75	0.01	0.08	0.13
G15	a ← b ↔ c	-1.75	0.01	0.08	0.13
G16	a ← b → c	-2.5	0	0.03	0.19

TABLE I: Breakdown of probability distribution and aggregate graded scores for each graph in a graph with 2 positive arguments and one negative

$$P(G|T) = \frac{P(T|G)P(G)}{\sum_{G \in S} P(T|G)P(G)} \tag{1}$$

In the formula above the variable $P(G)$ is the prior and represents our initial belief that graph $G \in S$ is the intended graph for view v before considering the observations. The term $P(T|G)$ is referred to as the likelihood and tells us that if we assume the intended graph for v is G what the probability (likelihood) of observing the observations T is. The multiple of the prior and the likelihood is proportional to the updated belief $P(G|T)$ which is referred to as the posterior. The denominator is called the normalising constant and ensures the resultant posterior distribution is a probability distribution; i.e. $\sum_{G \in S} P(G|T) = 1$. In the next sections we discuss our proposals for the prior and the likelihood functions.

III. Prior Distribution

In this section we summarise our previous proposal [1] for the prior distribution. Underpinning our approach is the assumption that the rating of a review can be used to identify the winning arguments in that review. Hence a high rating indicates the positive arguments are more likely win and vice versa. We formalise this intuition in threes parts.

The first component involves constructing a function, which we refer to as $\mathsf{Agg}(G)$, that scores each graph $G \in S$ in a graphs space, i.e $\mathsf{Agg} : S \to \mathbb{R}$. This score provides a measure of positivity based on an assessment of the graph's grounded extension and topology. This function enables us to score and rank the graphs in the graph space in terms of their positivity. The second component involves learning a function that maps between the rating of a review and this positivity domain. The third is to compute the mapped value for the rating in a review and compare this value to the positivity scores of the graphs in the graph space. These computed distances then serve as the basis for the identified probability distribution. For specific details of our implementation of the prior distribution please refer to our previous paper.

IV. Likelihood Distribution

In this section we detail our proposal for the likelihood function. This function allows us to update our prior distribution using data that captures the attacks between the arguments in the reviews we are analysing. Such data can be acquired for example using annotators where each annotator would be presented with a review and its constituent arguments and asked to identify the argument graph for that review based on their reading of the it. Using Bayesian terminology we refer these annotations as observations, T, and we recall that an observation $t \in T$ is a tuple (G, b) and $G = (A, R)$ is an argument graph.

In updating the probability distribution for a view we do not limit ourselves to using observations directly related to that view (i.e. annotations regarding that particular view). It may be the case that arguments present in that particular review are also present in others and so we would like to learn from annotations related to those reviews too. We therefore distinguish between these two types of observations.

The first is a **matching observation** in which the observation contains the exact same set of arguments as the ones in the view we are analysing. Formally we say given a view $v = (A_v, b_v)$ and an observation $t = ((A, R), b)$, t is a matching observation iff $A = A_v$. The second is a **similar observation** which contains a subset of arguments from the review we are considering. We say that an observation t is a similar observation iff it is not matching and the number of shared arguments is greater than or equal to 2; given that two arguments are required to have have attacks in the argument graph that can be used for learning, i.e $|\{A \cap A_v\}| \geq 2$.

Example 3. *Consider a view $(\{a, b, c\}, 10)$ and two observations $t_1 = ((\{a, b, c\}, \{\}), 10)$ and $t_2 = ((\{a, b\}, \{\}), 10)$. In this case t_1 is a matching observation and t_2 is a similar observation.*

Since a matching observation is contains an argument graph that can be found in the corresponding graph space for the view we are considering we can use the aggregate measure Agg as the basis for our likelihood function. However for a similar observation the aggregate measure cannot be used and so we will propose a modified likelihood function. In both cases we expect our resulting posterior distribution to adhere to certain rationality postulates which we discuss below.

(RPR) Rating Proportionality: Given two observations $t_1 = (G, b_1)$ and $t_2 = (G, b_2)$ and a view $(\mathsf{Args}(G), b)$ s.t $|b - b_1| > |b - b_2|$ then it should follow that $\sum_{G \in S} (P(G|t_1) - P(G))^2 > \sum_{G \in S} (P(G|t_2) - P(G))^2$. RPR states that for two observations that have the same graph but different ratings, the one that has the rating most similar to the rating of the view will impact the posterior most.

(APR) Argument proportionality: Given two observations $t_1 = (G_1, b_1)$ and $t_2 = (G_2, b_2)$ s.t $|\mathsf{Args}(G_1) \cap \mathsf{Args}(G)| \geq |\mathsf{Args}(G_2) \cap \mathsf{Args}(G)|$ and $b_1 = b_2$ it should follows that $\sum_{G \in S} (P(G|t_1) - P(G))^2 \geq \sum_{G \in S} (P(G|t_2) - P(G))^2$. APR states that for two observations, the observation that shares most arguments in common with the view will alter the posterior distribution most.

A. Updates with Matching Observations

We recall that the set of graphs that can be assigned to a view $v = (A_v, b_v)$ is $\mathsf{Space}(A^+ \cap A_v, A^- \cap A_v)$. For a matching observation $t = (G, b)$ we find that graph in the observation exists in this graph space, i.e. $G \in \mathsf{Space}(A^+ \cap A_v, A^- \cap A_v)$. What this means is that we can make use of the aggregate score defined in the previous section when determining the likelihood distribution.

Definition 4. *Given a view $v = (A_v, b_v)$ and a matching observation $t = (G_t, r_t)$ we define the relative* ***attack distance*** *between a graph $G \in S$ and observed graph G_t as $\mathsf{dist}_{\mathsf{Att}}(G, G_t) = 1$ iff $G_t = G$ and $\mathsf{dist}_{\mathsf{Att}}(G, G_t) = \frac{1}{2 + |\mathsf{Agg}(G) - \mathsf{Agg}(G_t)|}$ otherwise. We then say that the normalised distance is $\mathsf{normdist}_{\mathsf{Att}}(G, G_o) = \frac{\mathsf{dist}_{\mathsf{Att}}(G, G_o)}{\sum_{F \in S} \mathsf{dist}_{\mathsf{Att}}(F, G_o)}$.*

We now consider how the rating of an observation influences the resulting likelihood distribution. The basic idea is that observations that have ratings similar to the view should produce a likelihood distribution that influences the posterior more than those that are less similar. To illustrate consider a view $(A, 10)$ and two observations $t_1 = (G, 9)$ and $t_2 = (G, 1)$. In this case we would give greater priority to t_1 as the observation is very similar to the view's rating. In order to determine the degree to which an observation influences the posterior distribution we define a likelihood distribution which when used in Bayes' formulae has no effect whatsoever on the posterior, i.e. $P(G|T) = P(G)$. This **uninformative distribution** we know be the uniform distribution which we define as a $P(T|G) = \frac{1}{|S|}$.

We use the uninformative distribution as the likelihood distribution when difference in ratings between the observation and the view are maximally dissimilar. As the ratings become more and more similar we then gradually move towards just using the $\mathsf{normdist}$. We therefore say that given the rating difference is $\Delta b = \frac{b^+_{max} - b^-_{min} - |b - b_t|}{b^+_{max} - b^-_{min}}$. The update is then:

$$P(T|G) = \frac{1}{|S|} - \Delta b\left(\frac{1}{|S|} - \mathsf{normdist}_{Att}(G, G_t)\right) \quad (2)$$

Proposition 1. *The function we have defined for $P(T|G)$ in Equation 2 satisfies RPR.*

Proof. Let $t_1 = (G, b_1)$ and $t_2 = (G, b_2)$ be observations and $(\mathsf{Args}(G), b))$ a view s.t $|b_1 - r| > |b_2 - b|$. In both cases $\mathsf{normdist}_{Att}(G, G_t)$ is a constant term and the only differing term in the likelihood would be Δb. When Δb is high, i.e t_1, then the likelihood tends to $\frac{1}{|S|}$ which in turn means the posterior tends to $P(G)$ and thus $\sum_{G \in S} (P(G|t_2) - P(G))^2 > \sum_{G \in S} (P(G|t_1) - P(G))^2$ □

As can be seen the likelihood function is equal to $\mathsf{normdist}$ when $\Delta b = 1$, i.e. the view and the observations have the same rating, and equal to uniform distribution when $\Delta b = 0$. The likelihood distribution exhibits linear behaviour between these two extremes. We note that the likelihood function for matching observations trivially satisfies APR as all observations using this likelihood function always have the same set of arguments as the view in consideration.

B. Updates with Similar Observations

We recall that a similar observation is a relevant observation that does not have the exact same set of arguments as the view in consideration. In this case we are concerned with learning from the portion of arguments in the observation that overlap with the arguments in the view. In order to make this comparison we propose a similarity measure that compares the overlapping portion of the observed graph with the graphs in the graph space for the view based on topological structure as well as similarities in grounded extension. Measuring topological structure ensures that we are able to identify similarly structured graphs and measuring the similarity in grounded extension allows us to see how similar the conclusions are that both graphs reach. Further to this we enforce an additional measure which ensures that observations that have more arguments in common with the view's graph space contribute more to the update than observations that don't.

We begin by defining a similarity measure based on topological structure. We start by assessing the similarity between the four possible types of relations (**attack types**) between two arguments. For ease of notation we refer to the four possible attack types between any two arguments a, b given a set of attacks R as follows: we say $\mathsf{attackType}((a,b),R) = a \rightarrow b$ when $(a,b) \in R \,\&\, (b,a) \notin R$, $\mathsf{attackType}((a,b),R) = a \leftarrow b$ when $(b,a) \in R \,\&\, (a,b) \notin R$, $\mathsf{attackType}((a,b),R) = a \leftrightarrow b$ when $(a,b) \in R \,\&\, (b,a) \in R$, $\mathsf{attackType}((a,b),R) = a - b$ when $(a,b) \notin R \,\&\, (b,a) \notin R$. Where there is no confusion we also say that given an $attack \in \{a \leftarrow b, a \rightarrow b, a \leftrightarrow b, a - b\}$ and a graph $G = (A,R)$, $attack \in G$ iff $\mathsf{attackType}((a,b),R) = attack$. With these attack types we build a similarity measure to be able to measure the degree of similarity between two graphs in terms of their attacks.

Definition 5. *Let a, b be two arguments and let $X, Y \in \{a \leftarrow b, a \rightarrow b, a \leftrightarrow b, , a - b, \}$ be two attack types for a, b. Let $|\mathsf{arcs}(X)|$ represent the number of attacks in X. We define $\mathsf{dist_{att}}(X,Y)$ as:*

$$\mathsf{dist_{att}}(X,Y) = \begin{cases} 0 & \textit{if } X = Y \\ 1 & \textit{if } |\mathsf{arcs}(X)| \neq |\mathsf{arcs}(Y)| \\ 2 & \textit{if } |\mathsf{arcs}(X)| = |\mathsf{arcs}(Y)| \;\;\&\;\; X \neq Y \end{cases}$$

The measure captures the degree of dissimilarity between the attack types. We note that this measure is the same as the Hamming distance except in the case of $\mathsf{dist_{att}}(a \leftrightarrow b, a - b)$. The reason we give special treatment to these attack types is because we want to ensure that they are treated dissimilar to $\mathsf{dist_{att}}(a \leftrightarrow b, a \leftarrow b)$ or $\mathsf{dist_{att}}(a \leftrightarrow b, a \rightarrow b)$. This is because they have the same grounded extension and so in this sense are more similar.

To compute the overall similarity between two graphs we simply add the similarity scores for each pair of overlapping arguments as follows:

Definition 6. *Given graphs $G_1 = (A_1, R_1)$ and $G_2 = (A_2, R_2)$ the shared positive arguments are $A^+_{1,2} = A_1 \cap A_2 \cap A^+$ and the shared negative arguments are $A^-_{1,2} = A_1 \cap A_2 \cap A^-$. The set of possible conflicts as $C = A^+_{1,2} \times A^-_{1,2}$. For a given possible conflict $c \in C$ we define a distance $\mathsf{diff}(c, R_1, R_2) = 2 - \mathsf{dist_{att}}(\mathsf{attackType(c, R_1)}, \mathsf{attackType(c, R_2)})$. The degree of similarity between both graphs is then given by the function $\mathsf{similarity}(G_1, G_2) = \sum_{c \in C} \mathsf{diff}(c, R_1, R_2)$.*

Example 4. *Consider a graph $G_1 = (\{a, b\}, \{(a, b)\})$ and a larger graph $G_2 = (\{a, b, c\}, \{(a, b), (b, a), (b, c)\})$. The set of possible conflicts is $\{(a, b)\}$. From G_1 we find the $\mathsf{attackType}((a, b), \{(a, b)\}) = a \rightarrow b$ and from G_2 we find $\mathsf{attackType}((a, b), \{(a, b), (b, a), (b, c)\}) = a \leftrightarrow b$. The $\mathsf{similarity}(G_1, G_2) = 1$.*

Next we consider the degree to which two graphs are related in terms of their grounded extensions.

Definition 7. *Given two graphs $G_1 = (A_1, R_1)$ and $G_2 = (A_2, R_2)$ the overlap in the grounded extensions is $\mathsf{dist_{gr}}(G_1, G_2) = |(\mathsf{gr}(G_1) \setminus \mathsf{gr}(G_2)) \cup (\mathsf{gr}(G_2) \setminus \mathsf{gr}(G_1))|$ and that $\mathsf{dist}_{graph}(G_1, G_2) = \mathsf{similarity}(G_1, G_2) + \mathsf{dist_{gr}}(G_1, G_2)$. We then define the total distance between graphs as:*

$$\mathsf{dist_{total}}(G_1, G_2) = \frac{\mathsf{dist}_{graph}(G_1, G_2)}{\sum_{G \in S} \mathsf{dist}_{graph}(G, G_2)}$$

We finally want to ensure that the difference in rating between view and observation influences the final likelihood distribution. We also add an additional constraint which is that we want to make sure that an observation which is most similar to the graph space (proportion of overlapping arguments) contribute more to the update than those are less similar.

Definition 8. *Given an view (A, r) and a similar observation $t = (G_t, r_t)$. We say that $\Delta Args = \frac{|A \cap \mathsf{Args}(G)|}{\mathsf{max}(|A|, |\mathsf{Args}(G)|)}$. We then say that the likelihood function is this given by:*

$$P(T|G) = \frac{1}{|S|} - \Delta Args \Delta b \left(\frac{1}{|S|} - \mathsf{dist_{total}}(G, G_t)\right) \quad (3)$$

Example 5. *Let a, b be positive arguments, c be a negative argument, $v = (\{a, c\}, 10)$ a view and $t = ((\{a, b, c\}, (a, c)), 10)$. In this case $\Delta Args = 2/3$, $\Delta b = 1$.*

Proposition 2. *The function $P(T|G)$ defined in Equation 3 satisfies APR*

Proof. Let $t_1 = (G_1, b)$, $t_2 = (G_2, b)$ be observations where $G_1 = (A_1, R_1)$ and $G_2 = (A_2, R_2)$ s.t $G_1 \sqsubseteq G_2$ and let $v = (A, b')$ be a view. Δb is a constant and $\Delta Args$ is larger for t_1 than for t_2. Therefore $\sum_{G \in S} P(t_1|G)$ tends to $\frac{1}{|S|}$ as $\Delta Args$ decreases and likewise $\sum_{G \in S} P(t_1|G)$ tends to $\mathsf{dist}_{total}(G_1, G)$ as $\Delta Args$ grows. $\square$

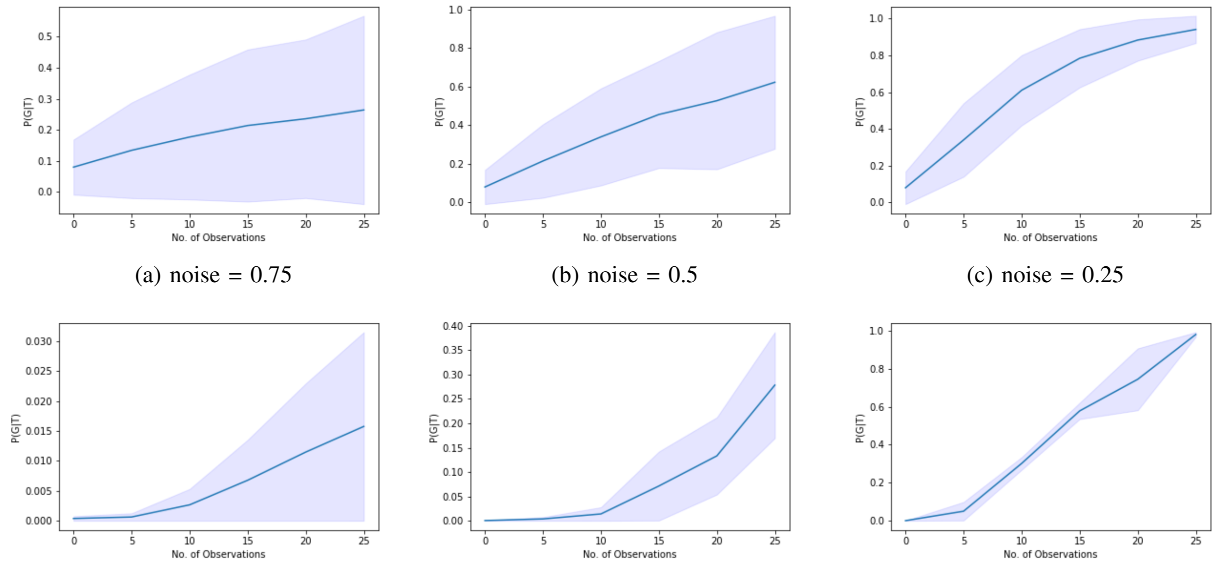

(a) noise = 0.75 (b) noise = 0.5 (c) noise = 0.25

(d) noise = 0.75 (e) noise = 0.5 (f) noise = 0.25

Fig. 2: Plots showing $P(G|T)$ against number of observations. The central blue lines Figures (a), (b), (c) show the average posterior probabilities for all $G \in \mathsf{Space}(\{a, b\}, \{c\})$ and (d), (e), (f) show them for all $G \in \mathsf{Space}(\{a, b, c\}, \{d, e\})$. The blue bands represent the 95% confidence intervals.

We note that our proposal can be adapted to settings in which ratings are not used, i.e where a view is simply a set of arguments and observations are argument graphs. In this case we can specify the uniform distribution as the prior distribution for the view and adapt our proposals for the likelihood function by making $\Delta r = 1$ as by doing so we remove the influence of the rating. Thus for a matching observation we would have a likelihood of $P(T|G) = \mathsf{normdist}_{Att}(G, G_t)$ and for a similar observation we would have $P(T|G) = \frac{1}{|S|} - \Delta Args(\frac{1}{|S|} - \mathsf{dist}_{total}(G, G_t))$.

V. Simulations

In this section we investigate the performance of our approach. Our primary motivation in these simulations is to demonstrate that our approach produces reasonable posterior probabilities under different circumstances. Firstly we expect that, for a given graph G, the posterior probability assigned $P(G|T)$, to be proportional to the frequency with which it appears in the observations. We refer to this as **noise** in observations w.r.t a graph of interest G and define it as follows.

$$\mathsf{noise}(G, T) = 1 - \frac{|\{t \in T | G \sqsubseteq t\}|}{|T|}$$

Where for two graph $G = (A, R)$ and $t = (A_t, R_t)$, $G \sqsubseteq t$ iff $A \subseteq A_t$ and $R = \{(a, b) \in R_t | a, b \in A\}$.

Example 6. *Let* $T = \{t_1, t_2\}$ *be a set of observations where* $t_1 = \{(\{a, b, c\}, (\{a, b, c\}, \{(a, b), (b, c)\}), 10)$ *and* $t_2 = \{(\{a, b\}, (\{a, b\}, \{(a, b)\}), 10)$. *We find that* $\mathsf{consistency}(a \to b, t) = 1$ *because* $a \to b$ *is in both observations.*

Another outcome we investigate is the effect that the number of observations (N) has on our model. In general we expect that as more observations are added the posterior probability should in general increase.

A. Synthetic Data

We created multiple synthetic datasets by using combinations of noise and number of observations. We started by specifying a set of positive arguments $\{a, b, c\}$ and negative arguments $\{d, e\}$. We experimented with 3 levels of noise $(0.25, 0.5, 0.75)$ and various numbers of observations in the range $5 \leq N \leq 25$. For each possible combination of noise and N we then created 150 synthetic datasets $(T_1, .., T_{150})$ for G where each graph in the set is randomly generated until we achieved the required level of noise and N. To simplify the simulations we the fixed the rating of every observation to 10. We then summarised the posterior probability for that combination of noise and N by computing an average posterior probability using all of the datasets i.e. $P(G|T) = \sum_{i=1}^{150} P(G|T_i)/150$.

B. Results

Figures 2a, 2b and 2c show how the posterior probability of all graphs in $\mathsf{Space}(\{a, b\}, \{d\})$ change with different levels of noise and dataset size. Instead of plotting the posterior probability for each individual graph in the graph space we plot the average of all the graphs in that graph space along with the 95% confidence interval bands denoted by the surrounding blue bands. Figures 2d, 2e and 2f show a similar picture but for a larger graph space of $\mathsf{Space}(\{a, b, c\}, \{d, e\})$. Note that

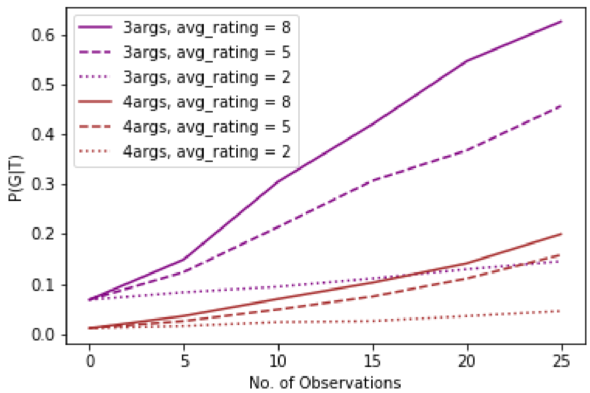

Fig. 3: Simulations showing the effect of ratings on the posterior probability

for all of the figures when the number of observation is 0 we report the prior probability for the graph.

As can be seen in all cases there is a general trend for the model to assign a higher posterior probability as more observations are added. Furthermore as noise decreases this learning is much faster. As can be seen clearly in Figures 2a, 2b, a larger starting prior noticeably influences the rate at which the posterior probability grows when noise is high. This is also noticeable in Figure 2d even though the graphs start with a prior probabilities that are all very small.

Figure 3 shows how using observations with different ratings effects the posterior probability. The purple lines show the average posterior probability for graphs in $\mathsf{Space}(\{a, b\}, \{d\})$ when the rating for the view is 10. Each of the three lines use observations with average ratings of 2, 5 and 8 respectively. The red lines show a similar picture but for graphs in $\mathsf{Space}(\{a, b\}, \{d, e\})$. What can be seen is that observations that use ratings very different to that of the view cause the posterior probability to grow slower than when they are similar; this demonstrates the effect of Δb on the posterior. For more detailed analysis of these simulations please refer to our supplementary results and analysis [1].

VI. Related Work

Learning in the context of abstract argumentation has been studied before. Proposals to learn argument frameworks from data in which the acceptability status of the arguments is known in advance have previously been made [5]. Many of these proposals also use probabilistic approaches [14] [7] [3] but ours differs fundamentally from all of these in that we do not assume access to such labels and instead we start from the assumption that we only have access to a proxy indicator (rating) through which we infer these acceptability statuses.

Proposals for the use of argumentation in analysing online reviews have been made in [4] which provides a method for extracting bipolar argument frameworks from film reviews where each reviewer provides labels the films as good or bad (binary choice). The approach differs from ours in that we provide a probabilistic output.

Lastly there have been proposals for instantiating argument graphs from online reviews [15]. Most related to our work is [9] which proposes a rule-based approach to build argument graphs from drug reviews. The proposal also assumes a relationship between ratings and acceptability status of arguments. However our approach is probabilistic and is able to learn from observations.

[1] https://github.com/robienoor/bayesianframeworklearningargumentation

VII. Discussion

In this paper we have proposed a method for identifying a prior distribution for a view in the absence of the agent's intended graph for that view and we have proposed a Bayesian inference framework that enables us to update the probability distribution for the view using data and therefore learn from the data.

We note that a current limitation of this approach exists in handling large numbers of arguments as this means generating and dealing with many arguments graphs in the graph space but we do intend to investigate combinatorial techniques to improve on this going forward [13]. Because positive/negative arguments are commonplace in the real world we have made this the focus in this paper however in future we intend to generalise our approach to handle k-partite graphs.

References

[1] Noor, Kawsar, and Anthony Hunter. "Analysing Product Reviews Using Probabilistic Argumentation". Proceedings of COMMA '20.

[2] Dung, Phan Minh. "On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games." Artificial intelligence 77.2 (1995): 321-357.

[3] Hunter, Anthony, and Kawsar Noor. "Aggregation of Perspectives Using the Constellations Approach to Probabilistic Argumentation." Proceedings of the AAAI '20.

[4] Cocarascu, Oana, Antonio Rago, and Francesca Toni. "Extracting dialogical explanations for review aggregations with argumentative dialogical agents." Proceedings of AAMAS '19.

[5] Riveret, Régis, and Guido Governatori. "On learning attacks in probabilistic abstract argumentation." Proceedings of AAMAS '16.

[6] Haenni, Rolf. Probabilistic argumentation." Journal of Applied Logic 7, no. 2 (2009): 155-176.

[7] Kido, Hiroyuki, and Keishi Okamoto. "A Bayesian Approach to Argument-Based Reasoning for Attack Estimation." Proceedings of IJCAI '17.

[8] Leite, Joao, and Joao Martins. "Social abstract argumentation." Proceedings of IJCAI '11.

[9] Noor, Kawsar, Anthony Hunter, and Astrid Mayer. "Analysis of medical arguments from patient experiences expressed on the social web." Proceedings of IEA/AIE '17.

[10] Li, Hengfei, Nir Oren, and Timothy J. Norman. "Probabilistic argumentation frameworks." Proceedings of TAFA '11.

[11] Hunter, Anthony. "Some foundations for probabilistic abstract argumentation." Proceedings of COMMA '12.

[12] Bonzon, Elise, et al. "A comparative study of ranking-based semantics for abstract argumentation." Proceeding of AAAI '16.

[13] Fazzinga, Bettina, Sergio Flesca, and Francesco Parisi. "Efficiently estimating the probability of extensions in abstract argumentation." Proceedings of SUM '13.

[14] Riveret, Régis et al. "A labelling framework for probabilistic argumentation". Annals of mathematics and artificial intelligence. Vol 83, No 1, Pages 21-71, 2018

[15] Schneider, Jodi, Tudor Groza, and Alexandre Passant. "A review of argumentation for the social semantic web." Semantic Web 4, no. 2 (2013): 159-218.