Opponent Modelling in Persuasion Dialogues

Christos Hadjinikolis, Yiannis Siantos, Sanjay Modgil, Elizabeth Black, Peter McBurney
King's College London, Department of informatics
London, UK
name.surname@kcl.ac.uk

Abstract

A strategy is used by a participant in a persuasion dialogue to select locutions most likely to achieve its objective of persuading its opponent. Such strategies often assume that the participant has a model of its opponents, which may be constructed on the basis of a participant's accumulated dialogue experience. However in most cases the fact that an agent's experience may encode additional information which if appropriately used could increase a strategy's efficiency, is neglected. In this work, we rely on an agent's experience to define a mechanism for augmenting an opponent model with information likely to be dialectally related to information already contained in it. Precise computation of this likelihood is exponential in the volume of related information. We thus describe and evaluate an approximate approach for computing these likelihoods based on Monte-Carlo simulation.

1 Introduction

Agents engaging in persuasion dialogues aim to convince their counterparts to accept propositions that the latter do not currently endorse. As one would expect, agents then strategise based on their assumptions about the beliefs of their counterparts (opponents). This is widely studied in terms of *opponent modelling*. Essentially, an opponent model (OM) consists of five basic components: an opponent's beliefs; abilities; preferences; objectives, and; strategy. Numerous researchers who deal with the *best response problem* in dialogues— optimally choosing which move to make — rely on opponent modelling for implementing, and employing strategies [Riveret *et al.*, 2007; 2008; Oren and Norman, 2010; Black and Atkinson, 2011]. The general idea is to rely on such a model, built from an agent's experience, for simulating the possible ways based in which a game may evolve. One may then rely on this simulation to optimally choose, from a number of options, the most suitable option with respect to one's goals. Most approaches exploit an agent's experience in a somewhat monolithic way, assuming that its opponent's beliefs can simply be modelled through collecting the distinct utterances that the latter puts forth in dialogues. Due to the simplicity of such modelling approaches, the formalisation of an OM is usually left implicit. In addition, such approaches disregard the fact that an agent's accumulated dialogue experience may encode additional information which could also be used for modelling, and so increasing the effectiveness of strategies that rely on OMs.

In this work, we rely on a simple framework for persuasion dialogue with abstract arguments. Participants maintain OMs which are constantly updated with content (arguments) obtained through new dialogues. We propose a mechanism which is used for augmenting such an OM with additional information that is likely to be associated with information already contained in the OM. In other words, we attempt to predict what else is *likely* to be believed by a particular agent, given: a) what we currently assume the latter believes, and; b) what others with similar beliefs know. To do this we rely on an agent's general history of dialogues, in which we monitor the times that certain opponent arguments (OAs) follow after certain others, thus utilising an agent's experience in a multifaceted way. For example, let two agents, Ag_1 and Ag_2, engage in a persuasion dialogue in order to decide where is the best place to have dinner:

- Ag_1:(A) We should go to the Massala Indian restaurant since a chef in today's newspaper recommended it.
- Ag_2:(B) A single chef's opinion is not trustworthy.
- Ag_1:(C) This one's is, as I have heard that he won the national best chef award this year.
- Ag_2:(D) Indian food is too oily and thus not healthy.
- Ag_1:(E) It's healthy, as it's made of natural foods and fats.

The above dialogue is essentially composed of two lines of dispute, $\{A \leftarrow B \leftarrow C\}$ and $\{A \leftarrow D \leftarrow E\}$. Assume then, that Ag_2 engages in a persuasion dialogue with another agent, Ag_3, on which is the best restaurant in town. Let us also assume that at some point in the dialogue Ag_3 cites the newspaper article, by asserting argument A in the game, as Ag_1 did in the previous dialogue. It is then reasonable for Ag_2 to expect that to some extent Ag_3 is *likely* to also be aware of the chef's qualifications (argument C). Intuitively, this expectation is based on a relationship between consecutive arguments in the same dispute lines of a dialogue. In this case, the "chef's proposition" (A) is defended against B's attack (i.e., 'supported') by "his qualifications" (C), suggesting some likelihood that awareness of the first implies awareness of the second. This is also the case for arguments A and E. However, assuming such a relationship between the chef's qualifications (C) and the argument on why Indian food is considered healthy (E), seems less intuitive, as these arguments belong in different dispute lines, and so do not support each other. We thus assume that two arguments can be related if they are found in the same dispute lines of a dialogue, where this relationship can be understood in terms of the notion of *support*, e.g., C supports A against B.

One can then rely on an agent's accumulated dialogue experience to define a graph in which links between OAs asserted in a series of dialogues, indicate support. A participant may then rely on this graph to augment an OM, by adding to it arguments which, according to the graph, are linked with arguments already contained in the OM, and which are thus likely to be known to that opponent. For quantifying this *likelihood* we rely on how often a certain argument follows after another in an agent's history of dialogues. While this is a simple approach, we believe that it is sufficient for increasing the effectiveness of an agent's strategising. In the future we intend to investigate more complex ways for quantifying this likelihood, accounting for contextual factors such as how common certain information is, or whether an agent is part of a certain group having access to shared information, etc.

Summarising, while an agent's dialogue experience can be utilised in many ways when modelling its opponents, in most cases this is neglected. We thus present an applicable method for updating and augmenting such a model utilising an agent's history of persuasion dialogues, making the following three contributions. In Section 3, we rely on an argumentative persuasion dialogue framework (formalised in Section 2) to define and present two mechanisms responsible for updating and augmenting an OM. Specifically; 1) Section 3.1 formalises a method for building a graph relating supporting arguments, based on an agent's experience; 2) Section 3.1 also provides a method for augmenting an agent's current beliefs about its opponents' beliefs based on these support relations, enabling an agent to additionally account for this information in its strategizing. Section 4 describes our third contribution; we define and analyse a Monte-Carlo simulation approach concerned with the augmentation process, which makes our approach tractable. We also prove convergence and provide supporting experimental results.

2 The Dialogue Framework

In this section we describe a framework for argumentation based persuasion dialogues in which participants exchange arguments constructed in a common language $\mathcal{L}$. Participants submit arguments that attack those of their opponent, and we assume that participants share an understanding of when one argument attacks another, based on the language dependent notion of conflict. The arguments $\mathcal{A}$ that are submitted, during the course of a dialogue are then assumed to be organised into a Dung framework, $AF = (\mathcal{A}, \mathcal{C})$, where $\mathcal{C}$ is the binary relation on $\mathcal{A}$, i.e. $\mathcal{C} \subseteq \mathcal{A} \times \mathcal{A}$. We assume participants take on the roles of proponent and opponent, where the former submits an initial argument X, whose claim is the topic of the dialogue. If X is a justified argument in the AF defined by the dialogue, then the proponent is said to have won the dialogue.

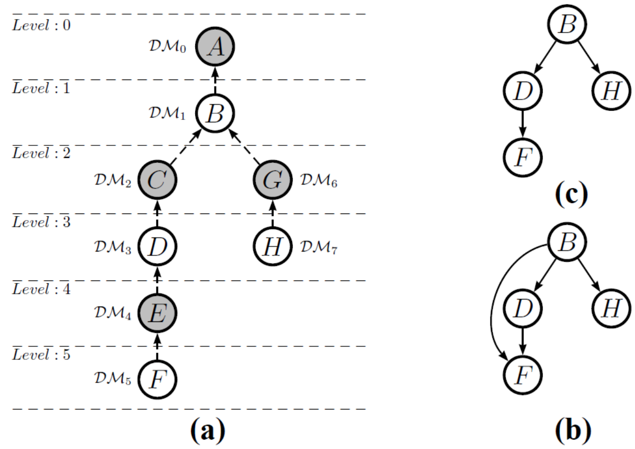

Figure 1: **(a)** A dialogue tree $\mathcal{T}$ where the grey and the white nodes concern Ag_1's respectively Ag_2's moves, **(b)** A 1-hop $\mathcal{RG}$ modelling approach **(c)** A 2-hop $\mathcal{RG}$ modelling approach.

We define a dialogue $\mathscr{D}$ as a sequence of dialogue moves $< \mathcal{DM}_0, \ldots, \mathcal{DM}_n >$, where the content of $\mathcal{DM}_0$ is the initial argument for the topic of the dialogue. We assume that the introduction of moves is contingent upon satisfying certain conditions, defined by a dialogue protocol, which concern: *turntaking*; *backtracking*; the *legality* of a move, and; the game's *termination rules*. Turntaking specifies the participant to move next and the number of dialogue moves she can make. We assume agents take alternate turns introducing a single move at a time. We assume multi-reply protocols which allow participants to backtrack and reply to previous moves of their interlocutors. Thus, a dialogue $\mathscr{D}$ can be represented as a tree $\mathcal{T}$:

Definition 1 *Let $\mathscr{D} =< \mathcal{DM}_0, \ldots, \mathcal{DM}_n >$ be a dialogue and $\mathcal{M} = \{\mathcal{DM}_0, \ldots, \mathcal{DM}_n\}$ the set of moves in $\mathscr{D}$. Then $\mathcal{T} = \{\mathcal{M}, \mathcal{E}\}$ is a* ***dialogue tree*** *with root node $\mathcal{DM}_0$, and arcs $\mathcal{E} \subseteq \mathcal{M} \times \mathcal{M}$, such that:*

- *for two moves $\mathcal{DM}_i$ & $\mathcal{DM}_j$, $(\mathcal{DM}_i, \mathcal{DM}_j) \in \mathcal{E}$ means that $\mathcal{DM}_j$ is $\mathcal{DM}_i$'s* target *($\mathcal{DM}_i$ replies to $\mathcal{DM}_j$)*
- *every move in $\mathcal{M}$ that is not the target of another move, is a* leaf *node*
- *Each distinct path from $\mathcal{DM}_0$ to a leaf node of $\mathcal{T}$, is referred to as a dispute.*

For a tree with m leaf nodes, $\Delta = \{d_1, \ldots, d_m\}$ is the set of all disputes in the tree. Each new dispute results from a backtracking move by either of the participants. An example is shown in Figure 1(a), where grey's move $\mathcal{DM}_6$ is used as an alternative reply against white's move $\mathcal{DM}_1$.

The legality of a dialogue move is regulated by explicit rules that account for the dialogical objective, a participant's role, and a participants commitments. The latter comprise a *commitment store* (CS) which is updated with the contents of the moves introduced by a participant during the course of the dialogue.

Definition 2 *The* ***commitment store*** *of an agent Ag participating in a dialogue $\mathscr{D} =< \mathcal{DM}_0, \ldots, \mathcal{DM}_n >$, at turn*

$t = 0 \dots n$, *is a set* $CS^t = \{A_0, \dots, A_k\}$, $k < n$, *containing the arguments introduced into the game by* Ag *up to turn* t, *such that* $CS^0 = \emptyset$, *and:*

$$CS^{t+1} = CS^t \cup \texttt{Content}(\mathcal{DM}_{t+1}) \quad (1)$$

where $\texttt{Content}(\mathcal{DM})$ *is the argument moved in* $\mathcal{DM}$.

Finally, we assume that each agent engages in dialogues in which its strategic selection of moves is based on what it believes its opponent believes. Accordingly each agent maintains a model of the beliefs of its potential opponents in terms of a set of arguments (as in [Oren and Norman, 2010]). Modelling an agent's goals is out of the scope of this work.

Definition 3 *Let* $\{Ag_1, \dots, Ag_\nu\}$ *be a set of agents. Then for* $i = 1 \dots \nu$, *the **belief base*** $\mathcal{KB}$ *of* Ag_i *is a tuple* $\mathcal{KB}_i =< \mathcal{A}_{(i,1)}, \dots, \mathcal{A}_{(i,\nu)} >$ *such that for* $j = 1 \dots \nu$, *each sub-base* $\mathcal{A}_{(i,j)} = \{A_1, \dots, A_k\}$, *where* $k \in \mathbb{N}$, *is a set representing an OM expressing what* Ag_i *believes is* Ag_j*'s known arguments; and where* $\mathcal{A}_{(i,i)}$ *represents* Ag_i*'s own beliefs.*

3 Modelling mechanisms

For modelling an agent's opponents' beliefs we rely on the arguments they put forth in a dialogue game. We, assume that agents believe what they utter, while. acknowledging that it is not possible to impose that agents are truthful through protocol restrictions or regulations [Wooldridge, 1999]. However, we believe our approach can readily be adapted to account for notions of trust and its use in argumentation [Y. Tang and S. Parsons, 2010]; a topic to be investigated in future work.

We begin by associating a *confidence* value c to the arguments of a sub-base $\mathcal{A}_{(i,j)}$. For an agent Ag_i this value expresses the probability of a certain argument in $\mathcal{A}_{(i,j)}$ being part of Ag_j's actual beliefs $\mathcal{A}_{(j,j)}$. To compute this value we differentiate between whether information is: a) gathered directly by Ag_i, on the basis of its opponent's updated CS, or; b) a result of Ag_i *augmenting* its current model of Ag_j. The latter involves incrementing an OM with the addition of arguments *likely* to also be known to Ag_i's opponent.

As noted in Section 1, intuitively we expect our opponents to be aware of arguments that are likely to follow in a current dialogue, given that they have appeared in previous dialogues, and relate to what we currently assume our opponents believe. We assume this likelihood to increase as the relation between the contents of an OM and the arguments external to the model becomes stronger. This may be due to the appearance of particular sequences of arguments in dialogues, which relate the two sets. For example, assume Ag_i (P) and Ag_j (O) engage in a persuasion dialogue, represented as the dialogue tree in Figure 1(a). In this case Ag_i and Ag_j introduce arguments $\{A, C, E, G\}$ and $\{B, D, F, H\}$ respectively. Assume then, that Ag_i engages in another persuasion dialogue with a different agent Ag_m who also counters Ag_i's A with B. It is then reasonable to assume that Ag_m is *likely* to be aware of arguments D, H or even F, given that D and H support B, and F supports D. If then Ag_m does indeed put forth arguments D, H and F in the game, then the likelihood of another agent knowing D, H and F, contingent that this other agent knows B, should increase.

For assigning a confidence value c to the elements of an $\mathcal{A}_{(i,j)}$, we assume every agent's set of arguments increase monotonically. This is compatible with the idea that the beliefs from which arguments are constructed are not revised upon incorporation of conflicting beliefs; rather the conflicts are resolved through evaluation of the justified arguments under acceptability semantics [Dung, 1995]. We therefore assume that the confidence value associated with arguments acquired by an an agent Ag from the commitment store of Ag's opponents (which we refer to as arguments 'directly collected by Ag') is equal to 1, which represents the highest level of confidence.

Definition 4 *Let* $\mathcal{A}_{(i,j)} \in \mathcal{KB}_i$. *Then* $\forall X \in \mathcal{A}_{(i,j)}$, c *is the **confidence level** that* Ag_i *associates with* X *(denoted by the tuple* $< X, c >$*) such that:*

$$c^{[0,1]} = \begin{cases} 1 & \textit{if } X \textit{ is directly collected by } Ag_i \\ Pr(X) & \textit{if } X \textit{ is part of an augmentation of } \mathcal{A}_{(i,j)} \end{cases}$$

where $Pr(X)$ *is the likelihood of* X *being known to* Ag_j.

Further details on how to determine $Pr(X)$ follow in Section 3.1. To define the mechanism for updating an agent's $\mathcal{A}_{(i,j)}$ we first need to define the notion of history:

Definition 5 *Let* $\{Ag_1, \dots, Ag_\nu\}$ *be a set of agents. For any two agents* Ag_i *and* Ag_j, $j \neq i$, $h_{(i,j)} = \{\mathscr{D}^1, \dots, \mathscr{D}^\mu\}$ *is* Ag_i*'s **history** of dialogues with* Ag_j. *Then* $\mathcal{H}_i = \bigcup_{j=1, j\neq i}^{\nu} h_{(i,j)}$ *is the set of all histories of* Ag_i *with each* $Ag_j, j \neq i$.

Commitment store updates to opponent models are then defined as follows:

Definition 6 *Let* $h_{(i,j)} = \{\mathscr{D}^1, \dots, \mathscr{D}^\mu\}$ *be the dialogue history of* Ag_i *and* Ag_j. *Given the current version of a sub-base* $\mathcal{A}^{\mu-1}_{(i,j)}$ *and* Ag_j*'s commitment store* CS_j *of the latest dialogue* $\mathscr{D}^\mu$, Ag_i *can **update** its sub-base as follows:*

$$\mathcal{A}^{\mu}_{(i,j)} = \mathcal{A}^{\mu-1}_{(i,j)} \cup CS_j$$

As explained above, in this case directly collected arguments are given a confidence value of 1.

3.1 Building a $\mathcal{R}G$ & augmenting the OM

Augmenting an $\mathcal{A}_{(i,j)}$ relies on an agent Ag_i's relationship graph ($\mathcal{R}G$) in which nodes are arguments asserted by Ag_i's opponents in $\mathcal{H}_i$. Nodes are related by weighted directed arcs which represent support relationships between arguments, while the weights represent the likelihood of these relationships. We assume $\mathcal{R}G_i$ is empty at the start (when $\mathcal{H}_i = \varnothing$) and incrementally updated with newly encountered opponent arguments (OAs) as Ag_i engages in dialogues. In the example shown in Figure 1(a), assuming the grey agent Ag_i is the modeller, then the OAs (the white's arguments B, D, F and H) can only appear in odd levels of the dialogue tree. Assigning arcs between these arguments relies on how and when an argument appears in a tree.

Essentially, we assume two OAs, X and Y, to be connected in a $\mathcal{R}G$ if they are found in the same path of the dialogue tree (i.e., dispute), and are of d-hop distance from each other (distance is measured disregarding the modeller's $\mathcal{DM}$s). For

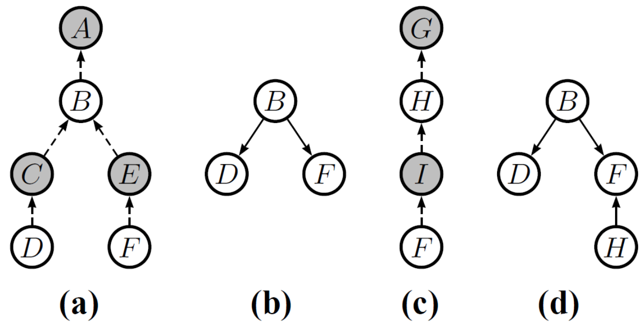

Figure 2: **(a)** A dialogue $\mathscr{D}^1$ between Ag_1 (grey) & Ag_2 (white), **(b)** The induced $\mathcal{R}G_1$, **(c)** A dialogue $\mathscr{D}^2$ between Ag_1 (grey) & Ag_3 (white), **(d)** The updated $\mathcal{R}G_1$

example, in Figure 1(a), B and D are of 1-hop distance from each other, while B and F are of 2-hop distance. Figures 1(b) and 1(c) illustrate two distinct $\mathcal{R}G$s induced by Figure 1(a)'s dialogue tree, for hop distances $d = 1$ and $d = 2$ respectively. Through modifying the d value one can strengthen or weaken the connectivity, between arguments in the same dispute, and correspondingly between arguments in the induced $\mathcal{R}G$. Of course, assigning a large d may raise *cognitive resources* issues, due to the large volume of information that needs to be stored. An example of incrementally building a $d = 1$ $\mathcal{R}G$ after two consecutive persuasion dialogues is illustrated in Figure 2.

Definition 7 *Let $\mathcal{A}^{\mathcal{H}}$ represent the arguments introduced by Ag's opponents in $\mathcal{H}$. Then $\mathcal{R}G$ is a directed graph $\mathcal{R}G = \{\mathcal{A}^{\mathcal{H}}, R\}$, where $R \subseteq \mathcal{A}^{\mathcal{H}} \times \mathcal{A}^{\mathcal{H}}$ is a set of weighted arcs representing support relationships. We write r_{AB} to denote the arc $(A, B) \in R$, and denote the arc's weight as w_{AB} obtained via a weighting function w, such that $w : R \to [0, 1]$.*

Whether two opponent's arguments in an agent's $\mathcal{R}G$ are connected is defined as follows:

Condition 1 *Let $\mathcal{T}$ be a dialogue tree $\mathcal{T}$ and d a hop distance. Then $r_{AB} \in R$ if there exists a distinct pair of opponent dialogue moves $\mathcal{DM}_1$ and $\mathcal{DM}_2$ with respective contents A and B, in the same dispute in $\mathcal{T}$, respectively appearing at levels i, j, $j > i$, where $\frac{j-i}{2} \leq d$;*

Lastly, for providing a weight value w_{AB} which will essentially represent the relationship likelihood of an argument A with an argument B, we rely on Definition 8, which is essentially a normalisation that allows us to compute a probability value $Pr(r_{AB}) = w_{AB}$ for arc r_{AB}. We simply count the number of agents that have used, in dialogues, argument A followed by B, and we put them against the total number of agents that have simply put forth argument A in dialogues.

Definition 8 *Given an Ag_i and its $\mathcal{R}G_i = \{\mathcal{A}^{\mathcal{H}_i}, R\}$, and two arguments A, B elements of $\mathcal{A}^{\mathcal{H}_i}$ then:*

$$\texttt{Occurances}(\mathcal{H}^i, A, B) = M_{AB}$$

is a function that returns a set $M_{AB} \subseteq Ags$ representing the set of agents that have put forth argument A followed by B in the same disputes and at a distance d in distinct dialogues in $\mathcal{H}^i$, satisfying Condition 1 such that $r_{AB} \in R$, then:

$$w_{AB} \equiv |M_{AB}|/|M_A| \qquad (2)$$

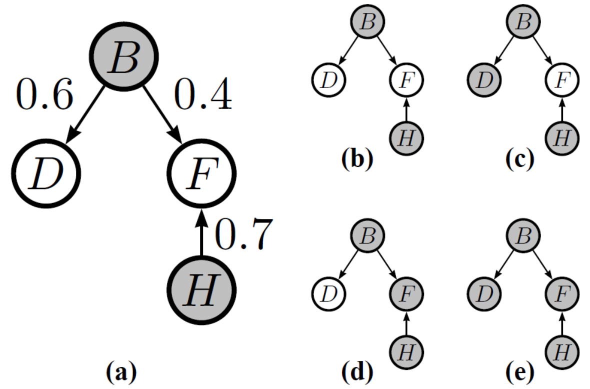

Figure 3: **(a)** $\mathcal{R}G_1$, **(b)**, **(c)**, **(d)** & **(e)** Possible augmentation $\mathcal{A}'_{\varnothing}$, $\mathcal{A}'_D$, $\mathcal{A}'_F$, & $\mathcal{A}'_{DF}$ respectively.

In the case where B is omitted then M_A will represent the set of agents that have simply put forth argument A in distinct dialogues, while it is evident that $|M_{AB}| \leq |M_A|$.

We should note that there is a problem with the aforementioned approach. Namely, if we consider the example shown in Figure 2, and based on Definition 8, then all the arcs in the induced $\mathcal{R}G$ will initially have a weight value of 1. It is apparent that this value does not represent the real likelihood which relates the arguments at either endpoint of the arc. In order to better approach the real likelihood a larger number of dialogues with numerous distinct participants need to be considered. This problem is better known as the *cold start problem* and is encountered in various other contexts as well [Lashkari *et al.*, 1994].

Having built a $\mathcal{R}G$ an agent Ag_i can then attempt to augment its OM ($\mathcal{A}_{(i,j)}$) of Ag_j by adding to it the possible arguments (nodes) that are of d-hop distance in $\mathcal{R}G$ from those contained in $\mathcal{A}_{(i,j)}$. In a trivial case, assume an $\mathcal{R}G_1$ induced by Ag_1 as it is illustrated in Figure 3(a), where for presentation convenience we assume that the weights on the arcs have received their values after numerous dialogue interactions. Let us assume that based on Ag_1's OM of Ag_4, Ag_1 believes that Ag_4 is aware of two arguments $\mathcal{A}_{(1,4)} = \{B, H\}$ (the grey nodes in Figure 3(a)). Hence, Ag_1 computes the *likelihood* of each of the possible augmentations $\mathcal{A}'_{(1,4)} \in P$ where $P = \{\mathcal{A}'_{\varnothing}, \mathcal{A}'_D, \mathcal{A}'_F, \mathcal{A}'_{DF}\}$, as those appear in Figures 3(b)(c)(d) and (e), and selects the one with the highest likelihood for augmenting $\mathcal{A}_{(1,4)}$ with additional contents. Computing each of these likelihoods is done as illustrated in the following example:

Example 1 *Assume we want to calculate the likelihood of augmentation $\mathcal{A}_{(1,4)} \to \mathcal{A}'_F$. In this simple example the likelihood of including belief F is:*

$$\begin{aligned} Pr(F) &= Pr(r_{HF} \cup r_{BF}) \\ &= Pr(r_{HF}) + Pr(r_{BF}) - Pr(r_{BF} \cap r_{HF}) \\ &= w_{HF} + w_{BF} - w_{BF} \cdot w_{HF} = 0.82 \end{aligned}$$

The probability of inducing $\mathcal{A}'_F$ therefore is the probability of including argument F and not including D which is:

$$Pr(\mathcal{A}'_F) = Pr(F)(1 - Pr(D)) = Pr(F)(1 - w_{BD}) = 0.328$$

Finally, $Pr(F)$ is also used to denote the confidence value c of argument F, as defined in Definition 4.

For providing the general formula for computing the likelihood of a possible augmentation we rely on basic graph theory notation with respect to a node X in a graph $\mathcal{R}G$, such as degree $d(X)$, neighbouring nodes $N(X)$ where $|N(X)| = d(X)$, and adjacent arcs $R(X)$. In addition, assuming a set of arguments $\mathcal{A}$, we define $N_{\mathcal{A}}$ such that $N_{\mathcal{A}} = \bigcup_{X \in \mathcal{A}} N(X) | \{Y \in N(X) : Y \notin \mathcal{A}\}$ and $R_{\mathcal{A}} = \bigcup_{X \in \mathcal{A}} R(X) | \{r_{XY} \in R(X) : Y \notin \mathcal{A}\}$. Essentially, set $N_{\mathcal{A}}$ represents the neighbours of the nodes in $\mathcal{A}$, and is formed from the union of the neighbours of every node X in $\mathcal{A}$, excluding those that are already in $\mathcal{A}$, while $R_{\mathcal{A}}$ is the adjacent arcs of the nodes in $\mathcal{A}$ and is equal to the adjacent arcs of every element X in $\mathcal{A}$, excluding those that connect with arguments already in $\mathcal{A}$. We note that for $\mathcal{A}_{(i,j)}$ it reasonably holds that $\mathcal{A}_{(i,j)} \subseteq \mathcal{A}^{\mathcal{H}_i}$, while for convenience we will hence refer to a $\mathcal{A}_{(i,j)}$ as $\mathcal{A}$ and to its augmentation as $\mathcal{A}'$. Given these, let $P = \{\mathcal{A}'_0, \mathcal{A}'_1, \ldots, \}$ be the set of all possible distinct augmentations of $\mathcal{A}$, then the number of all its possible distinct expansions with respect to neighbouring nodes that are of 1-hop distance from $\mathcal{A}$, is:

$$|P| = \sum_{k=0}^{|N_{\mathcal{A}}|} \binom{|N_{\mathcal{A}}|}{k} \tag{3}$$

The general formula for computing the likelihood of a possible augmentation $\mathcal{A}'$ with respect to the neighbouring nodes (arguments), i.e. for every $X \in N_{\mathcal{A}}$ of a set $\mathcal{A}$ is:

$$Pr(\mathcal{A}') = \prod_{X \in \mathcal{A}'} Pr(X) \prod_{X \notin \mathcal{A}'} (1 - Pr(X)) \tag{4}$$

$$Pr(X) = Pr(\bigcup_{r \in R_X^-} r) \tag{5}$$

where R_X^- (the in-bound arcs) is used to denote the adjacent arcs of X which are of the form (X, Y), i.e. Y is the arc target. Lastly, since the likelihood of each possible augmentation should define a distribution of likelihoods then it must also hold that $\sum_{\mathcal{A}' \in P} Pr(\mathcal{A}') = 1$.

4 The Monte-Carlo Simulation

A drawback of the proposed approach is that, calculating the probability of Equation 5 is of exponential complexity. This makes the approach practically intractable. However, drawing inspiration from the work of Li *et al.* [Li *et al.*, 2011] we rely on an approximate approach for computing these likelihoods based on a Monte-Carlo simulation. Essentially, the Monte-Carlo simulation attempts to estimate the argument likelihoods for a $\mathcal{R}G$ through sampling. In the case of a very large $\mathcal{R}G$ this approach is expected to be less expensive than exhaustively computing these likelihoods based on the inclusion-exclusion principle mentioned in Equation 5.

For the case of the 1-hop augmentation we know that the number of possible augmentations is exponential on the size of $N_{\mathcal{A}}$. It is possible to deduce a high likelihood augmentation in linear time, provided we know $Pr(Y)$ for $\forall Y \in N_{\mathcal{A}}$. Also, $Pr(Y)$ is calculated through Equation 5 using the inclusion-exclusion principle found in basic algorithm textbooks such as [Knuth, 1997]. However, this is of exponential complexity on the in-degree to calculate. We therefore proceed to sample for $Pr(Y)$, by describing a method to sample for high likelihood arguments which we will include in our 1-hop augmentation of a set $\mathcal{A}$. We will generally refer to the arguments in $\mathcal{A}$ as augmentation nodes. The method we describe is as follows:

- Assume an $\mathcal{R}G$= $\{\mathcal{A}^{\mathcal{H}}, R\}$ and a set $\mathcal{A}$
- We begin with a set of nodes $\mathcal{A}' = \mathcal{A}$
- For each argument $Y \in N_{\mathcal{A}}$ and if $\exists r_{XY} \in R$, where $X \in \mathcal{A}$, we accept r_{XY} with probability w_{XY}.
- If an arc r_{XY} is accepted, we then add Y to $\mathcal{A}'$
- At the end of the process $\mathcal{A}'$ contains a possible 1-hop augmentation of $\mathcal{A}$

We denote the probability of accepting an arc as $Pr(r_{XY}) = w_{XY}$. We also assume the events of accepting r_{ij} and r_{km}, where $r_{ij} \neq r_{km}$ to be independent but not mutually exclusive, therefore:

$$Pr(r_{ij} \cup r_{km}) = Pr(r_{ij}) + Pr(r_{km}) - Pr(r_{ij} \cap r_{km})$$

which means that the probability for a node to be added in $\mathcal{A}'$ follows Equation 5. Consequently the probability of obtaining a specific augmentation $\mathcal{A}'$ follows Equation 4, and therefore the described procedure essentially samples from the distribution of augmentations.

Assuming a sampling procedure which generates a number of n samples using the described method, then each node Y is included with probability equal to $Pr(Y)$. Thus after n independent identically distributed (*i.i.d.*) samples, a proportion of $nPr(Y)$ will contain argument Y. We define $k_Y = \sum_{i=1}^{n} I(Y \in \mathcal{A}'_i)$ where I is the indicator function taking the value of 1 if the predicate within is satisfied and $\mathcal{A}'_i$ is the set of nodes contained in the i-th sampled augmentation. In this case k_Y is used to denote the number of times we sampled Y. The expected number of augmentations samples which contain Y follows a binomial distribution. Thus the expected number of observations k_Y, of any given node Y after n tries is $\mathbf{E}\{k_Y\} = nPr(Y)$ and the variance is $\mathbf{Var}\{k_Y\} = nPr(Y)(1 - Pr(Y))$. This defines a multinomial distribution over the set of nodes. Due to the law of large numbers and the fact that each sample is *i.i.d.*, it holds that:

$$\mathbf{E}\{k_Y\} = nPr(Y) \quad \hat{Pr}(Y) = \frac{k_Y}{n}$$

where $\hat{Pr}(Y)$ is our estimate of $Pr(Y)$. The described procedure is a method to sample for the inclusion-exclusion probability. This is done since the exhaustive calculation of this probability is exponential in the number of arcs considered.

The generated random graphs are Poisson graphs, with edge probability set to $50/n$ (where n is the size of the graph). We selected this probability to ensure there exists a giant connected component and have a graph which is sufficiently dense, justifying the need to use a sampling method to infer the argument likelihoods rather than to directly measure them, but also sufficiently sparse, to be able to use it on a

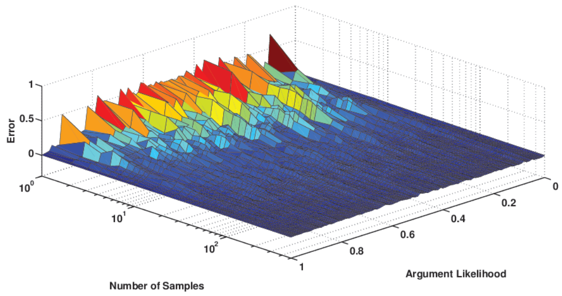

Figure 4: Error per argument likelihood over n samples

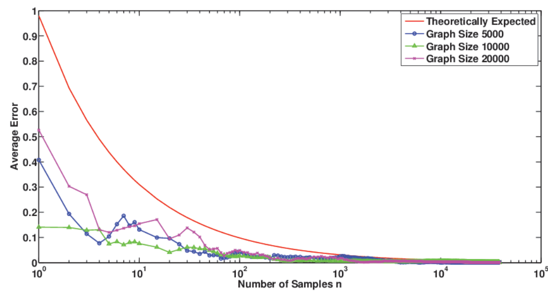

Figure 5: Average error over number of samples n

computer with limited capabilities. Due to the lack of benchmarks to compare with, we do not know whether a random graph better corresponds to a realistic argument graph. However, for the purposes of the Monte-Carlo simulation we believe that the graph structure is irrelevant for the purpose of argument likelihood estimation (assuming that the real argument graphs don't have a path-like or grid-like structure).

4.1 Sampling accuracy & Experimental results

The number of samples n required to achieve an accuracy ϵ with confidence δ is given through the following theorem[1]:

Theorem 1 *The Monte-Carlo approach to sample the probability distribution of the nodes in $N_{\mathcal{A}}$ is at least ϵ, close to convergence with probability at least δ after $n = z_\delta^2 \frac{1}{4\epsilon^2}$ samples.*

In essence, Theorem 1 gives us the expected upper bound of the error, which is expected to reach an accuracy equal to $\epsilon = 0.05$, with $\delta = 0.05$ confidence by taking $n \leq z_{0.975}^2 \frac{1}{4\epsilon^2} \approx 385$ samples, independently of the size of the $\mathcal{RG}$. Using the described algorithm we can estimate the likelihood of a given set of arguments (by dividing their observations by n) and infer a distribution of arguments. Based on these results we can augment our OM by choosing the set with the highest likelihood. Performing tests on randomly generated graphs of various sizes we have obtained the results of the average error which can be seen in Figure 5. We can see that the average error is upper bounded by the theoretically expected value. In practice the convergence is much quicker than what the theory suggests, resulting in error less than 0.1 after only 15 samples. Additionally in Figure 4 we can see the error per argument likelihood $Pr(Y)$.

[1] Space limitations preclude including proof of Theorem 1, which can be found in [C. Hadjinikolis, 2013].

5 Related Work & Conclusions

Specifically, in the context of dialogue games, Riveret *et al.* [Riveret *et al.*, 2007; 2008] model the possible knowledge of opponents in the form of arguments as we do. They however rely on the simplification that arguers are perfectly informed about all the arguments previously asserted by all their opponents in dialogues. Oren *et al.* [Oren and Norman, 2010] propose a variant of the min-max algorithm for strategising through relying on models which represent both an agent's knowledge, in the form of arguments, as well as their goals. However, nowhere in the aforementioned work is the problem of acquiring and maintaining, or augmenting an OM addressed. An exception, proposed by Black *et al.* [Black and Atkinson, 2011], concerns a mechanism that enables agents to model *preference* information about others—what is important to another agent—and then rely on this information for making proposals that are more likely to be agreeable. In their case the mechanism responsible for developing a model of an agent's preferences is explicitly provided, though they do not model agents' knowledge. A similar approach to our work has been proposed Rovatsos *et al.* [Rovatsos *et al.*, 2005] who explored how to learn stereotypical sequences of utterances in dialogues for deducing an opponent's strategy, though not relying on OMs. Finally, the work of Emele *et al.* [Emele *et al.*, 2011] is worth noting, as it is similar to ours in the sense that they also explore the development of an OM but based on what norms or expectations an opponent might have, and not on its general beliefs.

Through this work we have provided a general methodology for updating and augmenting an OM, based on an agent's experience obtained through dialogues. This methodology is based on two mechanisms respectively responsible for updating and augmenting an OM. In relation to the latter, we provided a method for building a graph between related arguments asserted by a modeller's opponents, referred to a $\mathcal{RG}$, and proposed an augmentation mechanism, enabling an agent to augment its current beliefs about its opponents beliefs by including additional information (arguments), that is of high likelihood to be related to what the opponent is currently assumed to know. Thus, we enabled an agent to also account in its strategizing for the possibility that additional information may also be known to its opponents.

Finally, we defined and analysed a Monte-Carlo simulation which enabled us to infer the likelihood of those additional arguments in a tractable and efficient way. We are aware that more investigation is needed with respect to relying on alternative contextual factors for quantifying the likelihood between elements in a $\mathcal{RG}$, such as the level of a participant's membership in a group, and this is something we intend to investigate in the future. We also intend to evaluate the effectiveness of our approach when employed in accordance with trust related semantics [Y. Tang and S. Parsons, 2010] through which one may define the trustworthiness of an agent's utterances against its actual beliefs.

References

[Black and Atkinson, 2011] E. Black and K. Atkinson. Choosing persuasive arguments for action. In *The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 3*, AAMAS–11, pages 905–912, 2011.

[C. Hadjinikolis, 2013] S. Modgil E. Black P. McBurney C. Hadjinikolis, Y. Siantos. Opponent modelling in persuasion dialogues. Technical Report, Department of Informatics, King's College London, 2013.

[Dung, 1995] P. M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2):321–357, 1995.

[Emele *et al.*, 2011] C. D. Emele, T. J. Norman, and S. Parsons. Argumentation strategies for plan resourcing. In *AAMAS*, pages 913–920, 2011.

[Knuth, 1997] D. E. Knuth. *Fundamental Algorithms*. Addison-Wesley Pub. Co, 1997.

[Lashkari *et al.*, 1994] Y. Lashkari, M. Metral, and P. Maes. Collaborative interface agents. In *In Proceedings of the Twelfth National Conference on Artificial Intelligence*, pages 444–449. AAAI Press, 1994.

[Li *et al.*, 2011] H. Li, N. Oren, and T. J. Norman. Probabilistic argumentation frameworks. In *TAFA*, pages 1–16, 2011.

[Oren and Norman, 2010] N. Oren and T. J. Norman. Arguing Using Opponent Models. In *Argumentation in Multi-Agent Systems*, volume 6057 of *Lecture Notes in Computer Science*, pages 160–174. 2010.

[Riveret *et al.*, 2007] R. Riveret, A. Rotolo, G. Sartor, H. Prakken, and B. Roth. Success chances in argument games: a probabilistic approach to legal disputes. In *Proceedings of the 20th annual conference on Legal Knowledge and Information Systems: JURIX*, pages 99–108, 2007.

[Riveret *et al.*, 2008] R. Riveret, H. Prakken, A. Rotolo, and G. Sartor. Heuristics in argumentation: A game-theoretical investigation. In *Proceedings of COMMA*, pages 324–335, 2008.

[Rovatsos *et al.*, 2005] M. Rovatsos, I. Rahwan, F. Fischer, and G. Weiss. G.: Adaptive strategies for practical argument-based negotiation. In *ETS 300 401, ETSI European Telecommunications Standards Institute*, 2005.

[Wooldridge, 1999] M. Wooldridge. Semantic issues in the verification of agent communication languages. *Autonomous Agents and Multi-Agent Systems*, 3, 1999.

[Y. Tang and S. Parsons, 2010] E. Sklar P. McBurney Y. Tang, K. Cai and S. Parsons. A system of argumentation for reasoning about trust. *In Proceedings of the 8th European Workshop on Multi-Agent Systems,*, 2010.