

Mechanisms for Opponent Modelling

Christos Hadjinikolis¹, Sanjay Modgil¹, Elizabeth Black¹, and Peter McBurney¹

1 Department of Informatics, King's College London, UK
name.surname@kcl.ac.uk

Abstract

In various competitive game contexts, gathering information about one's opponent and relying on it for planning a strategy has been the dominant approach for numerous researchers who deal with what in game theoretic terms is known as the *best response problem*. This approach is known as *opponent modelling*. The general idea is given a model of one's adversary to rely on it for simulating the possible ways based on which a game may evolve, so as to then choose out of a number of response options the most suitable in relation to one's goals. Similarly, many approaches concerned with strategising in the context of dialogue games rely on such models for implementing and employing strategies. In most cases though, the methodologies and the formal procedures based on which an opponent model may be built and updated receive little attention, as they are usually left implicit. In this paper we assume a general framework for argumentation-based persuasion dialogue, and we rely on a logical conception of arguments—based on the recent *ASPIC+* model for argumentation—to formally define a number of mechanisms based on which an opponent model may be built, updated, and augmented.

1998 ACM Subject Classification I.2. Computing Methodologies Artificial Intelligence

Keywords and phrases dialogue, strategies, argumentation, opponent model

Digital Object Identifier 10.4230/OASIS.ICCSW.2012.62

1 Introduction & Related work

Numerous researchers who deal with the best response problem rely on opponent modelling for implementing, employing and analysing strategies [1, 4, 3, 8, 10, 11, 12]. Essentially, an opponent model (OM) consists of four basic components: an opponent's knowledge; abilities; objectives, and; strategy. However, in most cases the methodologies and the formal procedures based on which such a model may be built and updated are often either left implicit, or are just concerned with particular components of the model.

Specifically, in the context of argumentation-based dialogue games, Riveret *et al.* [10, 11] model the possible knowledge of their opponents in the form of arguments, assuming that arguers are perfectly informed about all the arguments previously advanced by all other players. Their investigation concerns *games of perfect information*, and assumes that the participants' goals always comply with the dialogical objectives of the game, an assumption which, as McBurney *et al.* argue in [7], does not always hold. Oren *et al.* [8] present a generally complete approach through modelling both an agent's knowledge in the form of arguments as well as their goals, while in a similar sense to [4] they also allow for nested OMs. Additionally, they argue that given the knowledge about an opponent's goals it is also possible to indirectly model its strategy. However, nowhere in the aforementioned work is the problem of acquiring and maintaining an OM discussed. An interesting exception proposed by Black *et al.* [1] 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



© Christos Hadjinikolis, Sanjay Modgil, Elizabeth Black, and Peter McBurney;
licensed under Creative Commons License NC-ND
2012 Imperial College Computing Student Workshop (ICCSW'12).
Editor: Andrew V. Jones; pp. 62–68



OpenAccess Series in Informatics

OASIS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

ICCSW

proposals that are more likely to be agreeable. In their case the mechanism responsible for modelling an agent’s preferences is explicitly provided.

In this work, we attempt to formally address the problem of opponent modelling through providing two mechanisms concerned with how an OM may be built, updated and possibly augmented. The rest of the paper is structured as follows. In Section 2 we briefly present the basic components of an *ASPIC*⁺-based framework for persuasion dialogue. In Section 3 we present two mechanisms responsible for building, updating and augmenting an OM. Finally, in Section 4 we summarise and discuss future directions for our work.

2 An *ASPIC*⁺ framework for persuasion dialogues

In [9] Prakken describes an argumentation framework (*AF*) by assuming an unspecified *logical language* (\mathcal{L}) and by defining arguments as inference trees formed by applying *strict* (\mathcal{R}_s) or *defeasible* (\mathcal{R}_d) inference rules of the form $\varphi_1, \dots, \varphi_n \rightarrow \varphi$ respectively $\varphi_1, \dots, \varphi_n \Rightarrow \varphi$. To define attacks between arguments, minimal assumptions on \mathcal{L} are made; namely that certain well formed formulæ (wff) are a contrary or contradictory of certain other wff. Apart from this the framework is still abstract: it applies to any set of \mathcal{R}_s and \mathcal{R}_d and to any \mathcal{L} with a defined contrary relation. Arguments are then constructed with respect to a knowledge base that is assumed to contain three kinds of premises that are wff in \mathcal{L} . These premises are expressed through a set of three disjoint subsets $\mathcal{K} = \mathcal{K}_n \cup \mathcal{K}_p \cup \mathcal{K}_a$: \mathcal{K}_n is the (necessary) *axioms* (which cannot be attacked); \mathcal{K}_p is the *ordinary premises* (on which attacks succeed contingent upon satisfying certain preference criteria¹), and; \mathcal{K}_a is the *assumptions* (on which attacks are always successful, *cf.* assumptions in [2]). Thus an *ASPIC*⁺ argument in a set \mathcal{A} of arguments that may be constructed from the aforementioned logical components, may either be a single premise, or a chain of premises and rules that lead to a certain conclusion. Lastly, three kinds of *attack* are defined for arguments. $B \in \mathcal{A}$ can attack $A \in \mathcal{A}$ by attacking a premise or conclusion of A , or a defeasible inference step in A . Some kinds of attack succeed as *defeats* independently of preference criteria, whereas others succeed only if the attacked argument is not stronger than the attacking argument.

We assume a general framework for persuasion dialogue where the participants, a proponent (P) and an opponent (O), debate the truth of a claim φ through exchanging dialogue moves (\mathcal{DM} s) consisting of arguments, based on the attack relationship between them and on a set of protocol rules that regulate the dialogue process; i.e. a participant can introduce an argument into the game if it attacks another argument that was previously introduced into the game by its interlocutor. We assume that the participants share the same \mathcal{L} and the same contrary relation definition; i.e. there is agreement as to whether a given argument attacks another. In this respect, we define a dialogue \mathcal{D} as a sequence of dialogue moves $\langle \mathcal{DM}_0, \dots, \mathcal{DM}_n \rangle$, where the content of \mathcal{DM}_0 is an argument for φ , while we assume that the dialogue process is regulated by a multi-reply protocol. The latter means that *backtracking* is allowed, which implies that a participant can return to a previous point in the game and attack against a previous move of its interlocutor in a different way. Thus a dialogue may also be expressed in the form of a *dialogue tree* (\mathcal{T}), as the one illustrated in Figure 1(a), where each node is a \mathcal{DM} and each arc indicates a move’s target (a backtracking example is P ’s argument G introduced by \mathcal{DM}_6 in Figure 1(a)). Full details of the *ASPIC*⁺ framework as well as of the proposed dialogue framework can be found in [9] and [5] respectively.

¹ An important feature of the *ASPIC*⁺ framework is the employment of preference-orderings over defeasible rules and non-axiom premises which we do not take into account for the purpose of this paper.

We assume that the accumulated logical information introduced by a participant in \mathcal{D} is stored in a *commitment store* which we define as follows:

► **Definition 1.** Given a set of agents $Ag_s = \{Ag_1, \dots, Ag_\nu\}$ participating in $\mathcal{D} = \langle \mathcal{DM}_0, \dots, \mathcal{DM}_n \rangle$, then for any agent Ag_i we define its **commitment store** as a tuple $CSi^t = \langle \mathcal{K}i^t, \mathcal{R}i^t \rangle$, where $\mathcal{K}i^t$, and $\mathcal{R}i^t$, are respectively the premises and the rules moved into the game by Ag_i up to turn t , for $t = 0 \dots n$, such that $CSi^0 = \emptyset$, and CSi^{t+1} is obtained by augmenting CSi^t with the logical information provided by the dialogue move \mathcal{DM}_{t+1} .

Finally, we assume that each agent $Ag_i \in Ag_s$ can engage in dialogues in which its strategic selection of moves may be based on what Ag_i believes its interlocutor (in the set $Ag_{j \neq i}$) knows. Accordingly, and in a similar sense to the approach employed in [8], each Ag_i maintains a model of its possible opponent agents. In contrast with [8], the proposed model consists of the goals and knowledge other agents may use to construct arguments, rather than just the abstract arguments and their relations.

► **Definition 2.** Let $Ag_s = \{Ag_1, \dots, Ag_\nu\}$ be a set of agents. For $i = 1 \dots \nu$, the **knowledge base** \mathcal{KB} of Ag_i is a tuple $\mathcal{KB}_i = \langle S_{(i,1)}, \dots, S_{(i,\nu)} \rangle$ such that for $j = 1 \dots \nu$, each sub-base $S_{(i,j)} = \langle \mathcal{K}_{(i,j)}, \mathcal{R}_{(i,j)}, \mathcal{G}_{(i,j)} \rangle$ is an OM expressing what Ag_i believes is Ag_j 's premises ($\mathcal{K}_{(i,j)}$), rules ($\mathcal{R}_{(i,j)}$), and goals ($\mathcal{G}_{(i,j)}$), and where $S_{(i,i)}$ represents Ag_i 's own beliefs and goals.

3 Modelling mechanisms

We begin by associating a *confidence* value c to the logical components of the information sets found in a sub-base $S_{(i,j)}$. Essentially, for an agent Ag_i this value expresses the probability of a certain logical component in $S_{(i,j)}$ being part of Ag_j 's actual knowledge. For the computation of this value we differentiate between whether a particular information is: gathered directly by Ag_i , on the basis of its opponent's updated commitment store, or; a result of an *augmentation* attempt of Ag_i 's current model of Ag_j . The latter concerns an incrementation of a current OM with the addition of arguments that are *likely* to also be known to Ag_i 's opponent.

Intuitively, in real life we tend to assume that certain information, if known, is then likely to be related with other information, i.e. that there is some *relevance* between distinct pieces of information, which in our case may be translated as relevance between arguments. For example, assume that two agents, Ag_i (a proponent) and Ag_j (an opponent), engage in a dialogue as the one described in Figure 1(a), where Ag_i and Ag_j introduce arguments $\{A, C, E, G\}$ respectively $\{B, D, F, H\}$. Assume then, that Ag_i engages in another dialogue with the same root move A , but with a different agent Ag_m who also happens to counter A with argument B . It is then reasonable to assume that Ag_m is *likely* to also know of arguments D, H or even F . In this respect, the basic idea is, given an OM—which in essence describes a set of arguments known to one's opponent—and a mapping of a broader set of arguments with respect to a relevance factor, to then augment the OM by including arguments—and thus the logical elements that compose them—that have a high probability to also be known to that opponent, based on their relevance relationship with arguments already in the OM.

For assigning a confidence value c to the elements of an $S_{(i,j)}$, we will assume that every agent *retains* its own rules and premises without revision but relies on argumentation theory and semantics for resolving conflicts. In this respect, an agent's beliefs are formed based on deciding on the acceptability level of its arguments according to a number of different acceptability semantics. Thus in the face of new information *nothing is replaced or discarded*,

but instead certain arguments may simply cease to be or may become acceptable under different semantics. We will therefore assume that the confidence value of information acquired directly from the commitment store of one's interlocutor is equal to 1, which represents the highest level of confidence. However, this assumption must exclude information concerned with goals (\mathcal{G}) as those cannot be retained in the face of conflicts, i.e. it is not reasonable for an agent to be in pursue of conflicting goals at the same time. We leave the provision of a function for updating an opponent's goals to future work.

► **Definition 3.** Assume an $S_{(i,j)} \in \mathcal{KB}_i$, then for $Y \in \{\mathcal{K}_{(i,j)}, \mathcal{R}_{(i,j)}, \mathcal{G}_{(i,j)}\}$, X is a tuple $\langle x, c \rangle$ where $x \in Y$, and where c represents the **confidence level** of x such that:

$$c^{[0,1]} = \begin{cases} 1 & \text{if } x \text{ is directly collected by } Ag_i & (a) \\ Pr(x) & \text{if } x \text{ is part of an augmentation of } S_{(i,j)} & (b) \end{cases}$$

where $Pr(x)$ is the likelihood of x being also known to Ag_j , given its current OM ($S_{(i,j)}$).

► **Definition 4.** Let Ag_i and Ag_j be two agents in Ag_s such that $1 \leq i, j \leq \nu$ and $i \neq j$, and $h_{(i,j)} = \langle \mathcal{D}^1, \dots, \mathcal{D}^{\mu-1} \rangle$ be Ag_i 's history of dialogues with Ag_j . Then, given the current version of sub-base $S_{(i,j)}^{\mu-1} = \langle \mathcal{K}_{(i,j)}^{\mu-1}, \mathcal{R}_{(i,j)}^{\mu-1}, \mathcal{G}_{(i,j)}^{\mu-1} \rangle$ and the commitment store $CS_j = \langle \mathcal{K}_j, \mathcal{R}_j \rangle$ of the latest dialogue \mathcal{D}^μ , Ag_i can **update** its sub-base $S_{(i,j)}^\mu = \langle \mathcal{K}_{(i,j)}^\mu, \mathcal{R}_{(i,j)}^\mu, \mathcal{G}_{(i,j)}^\mu \rangle$ such that: (a) $\mathcal{K}_{(i,j)}^\mu = \mathcal{K}_{(i,j)}^{\mu-1} \cup \mathcal{K}_j$, and; (b) $\mathcal{R}_{(i,j)}^\mu = \mathcal{R}_{(i,j)}^{\mu-1} \cup \mathcal{R}_j$.

For augmenting a current OM, we rely on a *relevance graph*.

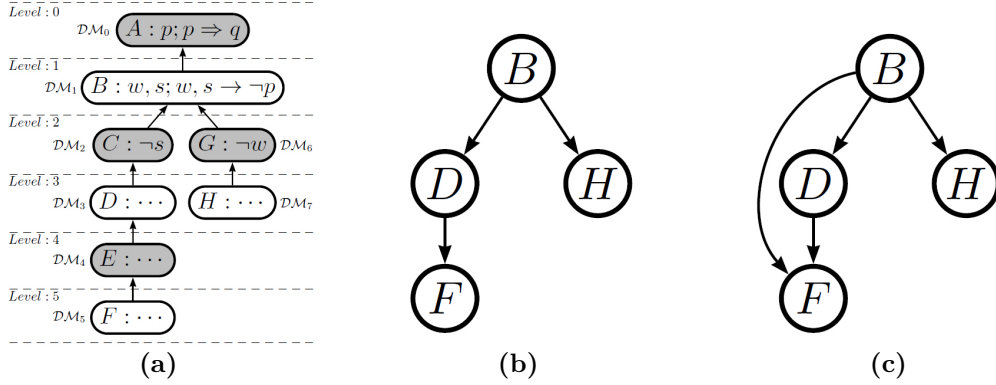
► **Definition 5.** For an agent Ag_i , let $\mathcal{H}_i = \{h_{(i,1)}, \dots, h_{(i,\nu)}\}$ be the set of all its histories, then an **abstract relevance graph (ARG)** is a weighted directed graph $\mathcal{G} = \{V, R\}$, where V is a set that consists of all arguments $A^{\mathcal{H}}$ encountered by Ag_i in \mathcal{H}_i , and where R is a set of weighted arcs, each of them indicating a relevance relationship between two arguments in \mathcal{G} , based on a weight function w , such that $w : R \rightarrow [0, 1]$.

3.1 Building a relevance graph

We assume an ARG to be incrementally built as an agent Ag_i engages in numerous dialogues, being empty at the beginning, and constantly updated with newly encountered opponent arguments (OAs). Notice that OAs appear only in the odd levels of a tree (Figure 1(a)). For assigning arcs between these arguments one may rely on how and when an argument appears in a dialogue tree. Specifically, we rely on the following condition:

► **Condition 1.** Given a dialogue tree \mathcal{T} , then for any argument A that appears in level i , and any argument B that appears in level j , for i and j being odd numbers and $j \geq i$, if $\frac{j-i}{2} \leq n$, and there exists a path between A and B in \mathcal{T} , there is an arc from A to B in \mathcal{G} .

Figures 1(b) and 1(c) illustrate two distinct ARGs induced from the dialogue tree of Figure 1(a), for $n = 1$ and $n = 2$ respectively. Intuitively, this modelling approach simply reflects the implied relationship that consecutive OAs have in a single branch of a tree. Through modifying the n value one can strengthen or weaken the connectivity, and so the relationship, between arguments in the induced ARG. However, one may choose to deviate from this particular modelling approach, adopting a different one so as to reflect a different kind of implied relationship between arguments. Lastly, for a pair of arguments $\{A, B\}$ connected with an arc r , let $w(r_{AB})$ be the weight value of an arc r which extends from argument A to argument B , we assume $w(r_{AB})$ to be equal to the number of agents N_{AB} that have moved A followed by B in a dialogue game, thus satisfying the relevance condition for $n = 1$, against the total number of agents $|Ag_s|$ minus agent Ag_i , i.e.



■ **Figure 1** (a) A dialogue tree \mathcal{T} where the grey and the white nodes concern P 's respectively O 's moves, (b) A 1-hop ARG modelling approach (c) A 2-hop ARG modelling approach.


$w(r_{AB}) = N_{AB}/(|Ags| - 1)$. Dividing N_{AB} with $(|Ags| - 1)$ is necessary for normalising the arcs' weight values into probabilities.

3.2 Relevance augmentation

Given an ARG an agent Ag_i can then attempt to augment its OM of Ag_j (i.e. its $S_{(i,j)}$) by adding to it the logical information comprised in the arguments (nodes) that are of 1-hop distance in \mathcal{G} from those that can be constructed from $S_{(i,j)}$. In a trivial case, let $Ags = \{Ag_1, Ag_2, Ag_3, Ag_4\}$, and \mathcal{G} be an ARG induced by Ag_1 based on dialogues \mathcal{D}^1 and \mathcal{D}^2 for $n = 1$, as it appears in Figure 2(c). Let $S_{(1,4)}$ be Ag_1 's OM of Ag_4 such that Ag_1 believes that Ag_4 can construct two arguments $\mathcal{A} = \{B, H\}$. Thus Ag_1 's OM of Ag_4 can be expressed as a sub-graph $G_{\mathcal{A}} = \{\mathcal{A}, \emptyset\}$ of \mathcal{G} (the yellow nodes in Figure 2(c)). Hence, Ag_1 computes the *likelihood* of each of the possible augmentations $\mathcal{A}' \in P$ of \mathcal{A} as those appear in set $P = \{\mathcal{A}'_{\emptyset}, \mathcal{A}'_D, \mathcal{A}'_F, \mathcal{A}'_{DF}\}$ (Figure 2(c)), and selects the one with the highest likelihood for augmenting $S_{(1,4)}$. Given that the instantiation of an augmented OM relies on the arcs' weights of \mathcal{G} , we have to provide an arc-centric formula for computing this probability, as there are multiple ways based on which a particular augmentation may be induced. In other words and in graph theoretic terms, a possible augmentation of an OM is interpreted as a possible *graph expansion*. Thus a certain \mathcal{A}' may be induced as a result of numerous possible expansions of $G_{\mathcal{A}}$, each containing *different arcs* while having the same set of arguments.

For example, assume we want to calculate the likelihood of augmentation $\mathcal{A}'_F = \{B, H, F\}$. Let $S(\mathcal{A}'_F) = \{G_1, G_2, G_3\}$ be the possible expansions of $G_{\mathcal{A}}$ which induce \mathcal{A}'_F by including in $G_{\mathcal{A}}$: either only r_{BF} creating G_1 ; either only r_{HF} creating G_2 , or; only r_{BF} and r_{HF} creating G_3 . Hence, the likelihood of \mathcal{A}'_F is: $Pr(\mathcal{A}'_F) = Pr(G_1) + Pr(G_2) + Pr(G_3) \Leftrightarrow Pr(\mathcal{A}'_F) = w(r_{BF}) \cdot (1 - w(r_{BD})) \cdot (1 - w(r_{HF})) + w(r_{HF}) \cdot (1 - w(r_{BF})) \cdot (1 - w(r_{BD})) + w(r_{BF}) \cdot w(r_{HF}) \cdot (1 - w(r_{BD})) \Leftrightarrow Pr(\mathcal{A}'_F) = 0.35937$. Finally, the confidence value c of the newly included information in $S_{(1,4)}$ is assigned a value equal to the likelihood of the chosen augmentation as defined by Definition 3(b), i.e. $Pr(x) = Pr(\mathcal{A}')$.

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 A in a graph \mathcal{G} , such as degree $d(A)$, adjacent vertices $N(A)$ where $|N(A)| = d(A)$, adjacent arcs $R(A)$, and arc weights $w(r)$. We additionally define N_S for a set of arguments S such that $N_S = \bigcup_{A \in S} N(A) \setminus \{X \in N(A) : X \notin S\}$, and $R_S = \bigcup_{A \in S} R(A) \setminus \{r_{AB} \in R(A) : B \notin S\}$. Additionally, let \mathcal{A} be the



■ **Figure 2** (a) A dialogue \mathcal{D}^1 between Ag_1 & Ag_2 (b) A dialogue \mathcal{D}^2 between Ag_1 & Ag_3 (c) An ARG induced by Ag_1 from \mathcal{D}^1 & \mathcal{D}^2 , and the image of Ag_1 's OM of Ag_4 on it (the yellow nodes B & H).

set of all arguments that may be induced from a single sub-base $S_{(i,j)}$, then given that an ARG is essentially built from a number of OMs, then it must hold that $\mathcal{A} \subseteq \mathcal{A}^{\mathcal{H}}$. Provided \mathcal{A} , we assume \mathcal{A}' to be an augmentation of \mathcal{A} based on \mathcal{G} , such that $\mathcal{A}' = \mathcal{A} \cup S$ for $S \subseteq N_{\mathcal{A}}$. In this sense, we assume $G_{\mathcal{A}} = \{\mathcal{A}, \emptyset\}$ to be a sub-graph of \mathcal{G} representing an image of an agent's Ag_i current OM of another agent Ag_j in \mathcal{G} , while we also assume $G_{\mathcal{A}'} = \{\mathcal{A}', R_i\}$ to be a possible expansion of $G_{\mathcal{A}}$, where $R_i \subseteq R_{\mathcal{A}}$. Given these, let $P = \{\mathcal{A}'_0, \mathcal{A}'_1, \dots, \mathcal{A}'_{\mu-1}\}$ be the set of all possible distinct augmentations of \mathcal{A} , then the number of all possible distinct expansions of $G_{\mathcal{A}}$ with respect to neighbouring nodes that are of 1-hope distance from it, is:

$$\mu = |P| = \sum_{i=0}^{|N_{\mathcal{A}}|} \binom{|N_{\mathcal{A}}|}{i} \quad (1)$$

Furthermore, let $S(\mathcal{A}') = \{G_{\mathcal{A}'}^1, \dots, G_{\mathcal{A}'}^n\}$ be a set of graphs containing all expanded graphs that have the same set of arguments \mathcal{A}' such that $G_{\mathcal{A}'} = \{\mathcal{A}', R_j\}$ for $R_j \subseteq R_{\mathcal{A}}$, then the general formula for computing the likelihood of a possible augmentation is:

$$Pr(\mathcal{A}') = \sum_{G_{\mathcal{A}'}^{i=1} \in S(\mathcal{A}')} \left(\prod_{r \in R_j} w(r) \cdot \prod_{r \in R_{\mathcal{A}}/R_j} (1 - w(r)) \right) \quad (2)$$

Finally, since the likelihood of each possible augmentation should define a distribution of likelihoods then it must hold that:

$$\sum_{\mathcal{A}'_{i=0}}^{\mu-1} Pr(\mathcal{A}'_i) = 1 \quad (3)$$

4 Conclusions & Future direction

In this work we have addressed the problem of building, updating and augmenting an OM in argumentation-based dialogues. We relied on a logical conception of arguments based on the recent *ASPIC⁺* model for argumentation, and provided two modelling mechanisms: an update mechanism, and; an augmentation mechanism.

We have particularly focused on the latter, which relies on the relevance between information and attempts an augmentation of a current OM through the addition of information.

The latter is based on computing the likelihoods of a set of possible augmentations and choosing the one with the highest value. A drawback of the proposed approach is that, given a \mathcal{G} , all possible augmentations of a set \mathcal{A} is equal to $|P| = 2^{|\mathcal{N}(\mathcal{A})|}$ (each adjacent argument is either in or out of the augmentation), where $\mathcal{N}_{\mathcal{A}}$ is the adjacent arguments of \mathcal{A} (its neighbours), which implies that the complexity of computing the likelihoods of all possible augmentations of \mathcal{A} increases exponentially as the number of the 1-hop neighbours of \mathcal{A} increases. This makes the approach practically intractable.

However, drawing inspiration from the work of Li *et al.* [6] we intend to rely on an approximate approach for computing these likelihoods based on a Monte-Carlo simulation. Therefore our immediate future direction is to formally describe the exact simulation process for the proposed augmentation method. Additionally, we also intend to evaluate our approach through experimenting with software agents that engage in dialogue disputes. Particularly, we will compare the success rate of agents that rely on OMs and a relevance augmentation mechanism to agents who rely on simple opponent modelling and to agents who do not rely on opponent modelling at all.

References

- 1 E. Black and K. Atkinson. Choosing persuasive arguments for action. In *10th International Conference on Autonomous Agents and Multi-Agent Systems*, 2011.
- 2 Andrei Bondarenko, Phan Minh Dung, Robert Kowalski, and Francesca Toni. An abstract, argumentation-theoretic approach to default reasoning. *Artificial Intelligence*, 93(1–2):63–101, 1997.
- 3 D. Carmel and S. Markovitch. Model-based learning of interaction strategies in multi-agent systems. *Journal of Experimental and Theoretical Artificial Intelligence*, 10(3):309–332, 1998.
- 4 David Carmel and Shaul Markovitch. Incorporating opponent models into adversary search. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence*, pages 120–125, Portland, Oregon, 1996.
- 5 C. Hadjinikolis, S. Modgil, E. Black, P. McBurney, and M. Luck. Investigating Strategic Considerations in Persuasion Dialogue Games. In *STAIRS 2012 - Proceedings of the Sixth Starting AI Researchers' Symposium*, volume 241 of *Frontiers in Artificial Intelligence and Applications*, pages 137–148. IOS Press, 2012.
- 6 Hengfei Li, Nir Oren, and Timothy J. Norman. Probabilistic argumentation frameworks. In *TAFAs*, pages 1–16, 2011.
- 7 P. McBurney and S. Parsons. Games That Agents Play: A Formal Framework for Dialogues between Autonomous Agents. *Journal of Logic, Language and Information*, 11(3):315–334, 2002.
- 8 N. Oren and T. Norman. Arguing Using Opponent Models. In *Argumentation in Multi-Agent Systems*, volume 6057 of *Lecture Notes in Computer Science*, pages 160–174. 2010.
- 9 H. Prakken. An abstract framework for argumentation with structured arguments. *Argument and Computation*, 1(2):93–124, 2010.
- 10 R. Riveret, H. Prakken, A. Rotolo, and G. Sartor. Heuristics in argumentation: A game-theoretical investigation. In *Proceedings of COMMA*, pages 324–335, 2008.
- 11 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.
- 12 William E. Walsh, Rajarshi Das, Gerald Tesauro, and Jeffrey O. Kephart. Analyzing complex strategic interactions in multi-agent systems. In *AAAI-02 Workshop on Game-Theoretic and Decision-Theoretic Agents*, Edmonton, 2002.