Opponent Models with Uncertainty for Strategic Argumentation

Tjitze Rienstra
Individual and Collective Reasoning
University of Luxembourg, Luxembourg
tjitze.rienstra@uni.lu

Matthias Thimm
Institute for Web Science and Technologies
Universität Koblenz-Landau, Germany
thimm@uni-koblenz.de

Nir Oren
Department of Computing Science
University of Aberdeen, Scotland, UK
n.oren@abdn.ac.uk

Abstract

This paper deals with the issue of strategic argumentation in the setting of Dung-style abstract argumentation theory. Such reasoning takes place through the use of opponent models—recursive representations of an agent's knowledge and beliefs regarding the opponent's knowledge. Using such models, we present three approaches to reasoning. The first directly utilises the opponent model to identify the best move to advance in a dialogue. The second extends our basic approach through the use of quantitative uncertainty over the opponent's model. The final extension introduces *virtual arguments* into the opponent's reasoning process. Such arguments are unknown to the agent, but presumed to exist and interact with known arguments. They are therefore used to add a primitive notion of *risk* to the agent's reasoning. We have implemented our models and we have performed an empirical analysis that shows that this added expressivity improves the performance of an agent in a dialogue.

1 Introduction

Argumentation systems offer a natural, easily understood representation of non-monotonic reasoning, and have been applied to a variety of problem domains including planning and practical reasoning [Toniolo *et al.*, 2011] and legal reasoning [Grabmair and Ashley, 2010]. Critically, many of these domains are adversarial, requiring an agent to identify and advance some set of arguments which are most likely to enable it to achieve its goals. In order to do so, the agent employs a strategy, typically in the form of a heuristic, which selects appropriate arguments given some contextual knowledge.

We describe one such strategy, and examine some of its properties. Our strategy assumes that an agent is not only aware of the arguments that it is permitted to advance, as well as what has already been stated, but that it also has a belief regarding its opponent's knowledge, and that this relationship is recursive unto some depth (i. e. an agent a believes some arguments, and believes that b knows some arguments, as well as believing that b believes that a knows some arguments, and so on). While [Oren and Norman, 2009] have previously examined such a strategy, we extend and improve their work along several dimensions.

First, [Oren and Norman, 2009] assume that an agent holds only a single opponent model. However, uncertainty plays a crucial role in strategies, and we capture this uncertainty, associating different opponent models with different likelihoods. Second, agents are often unaware of all arguments in a domain, and we allow an agent to hold an opponent model containing arguments it itself is not aware of through the introduction of *virtual arguments*. Finally, we consider how an agent should update its knowledge and beliefs while taking part in a dialogue.

In [Prakken and Sartor, 2002] an influential four layered view of an argumentation system is described. The first two levels, consisting of the logical and dialectic layers, specify the content of an argument, and how arguments interact with each other. In our work, these layers are encapsulated within an abstract argumentation framework [Dung, 1995], which we summarise in Sec. 2. We encapsulate Prakken's *procedural layer*, which specifies how agents exchange arguments (via dialogue) via a general *discourse model*. This discourse model, described in Sec. 3, assumes only that agents take alternating turns in making moves, and further constrains the dialogue by limiting what moves can be made through a *legal moves* function. Section 4 then describes the *agent model*, associating a utility with specific arguments, and allowing for different types of belief states. This agent model captures Prakken's heuristic layer through the specification of agent strategy. We present three instances of the agent model, starting from the one described in [Oren and Norman, 2009], and repeatedly adding further layers of expressivity. Section 5 describes how an agent's beliefs should be updated as the dialogue progresses, following which we compare and empirically evaluate the different models in Sec. 6. Section 7 discusses related work, and we conclude in Sec. 8.

2 Formal preliminaries

In abstract argumentation theory, knowledge is represented by an abstract *argumentation framework* (or AF, in short), which is a set of arguments with an attack relation, cf. [Dung, 1995].

Definition 1 *An* AF *is a pair* (A, R) *where* A *is the set of* arguments *and* $R \subseteq A \times A$ *is the* attack relation.

The goal is to select sets of arguments, called *extensions*, that represent rational points of view on the acceptability of the arguments of the AF. The first condition for an extension to be rational is that it is *conflict-free*. Furthermore, if an argument is a member of an extension, it is assumed that it is *defended* by the extension. Formally:

Definition 2 *Given an AF* $F = (A, R)$*, an* extension *is a set* $E \subseteq A$. E *is said to be* conflict-free *iff* $\nexists x, y \in E, (x, y) \in R$. *Given an argument* $x \in A$, E *is said to* defend x *iff* $\forall y \in A$ *s.t.* $(y, x) \in R$, $\exists z \in E$ *s.t.* $(z, y) \in R$. *We define* $D_{(A,R)}(E)$ *by* $D_{(A,R)}(E) = \{x \in A \mid E \text{ defends } x\}$.

Using the notions of conflict-freeness and defense, we can define a number of argumentation semantics, each embodying a particular rationality criterion.

Definition 3 *Let* $F = (A, R)$ *be and* $E \subseteq A$ *conflict-free extension.* E *qualifies as:*

- ***admissible*** *iff* $E \subseteq D_{(A,R)}(E)$,
- ***complete*** *iff* $E = D_{(A,R)}(E)$,
- ***grounded*** *iff* E *is minimal (w.r.t. set inclusion) among the set of complete extensions of* F.

For the intuition behind the different semantics we refer the reader to [Dung, 1995].

3 The Discourse Model

The *discourse model* provides a way to specify the complete setting in which two agents (proponent $\mathcal{P}$ and opponent $\mathcal{O}$) engage in a certain type of dialogue. Recalling the four-layered model mentioned in the introduction, we first need a logical and dialectical layer. Here we use abstract argumentation theory, as presented in the previous section, and leave the logical content of arguments unspecified. We assume that there is a *universal AF* $(\mathcal{A}, \mathcal{R})$ which contains all arguments relevant to a particular discourse.

An agent $Ag \in \{\mathcal{P}, \mathcal{O}\}$ has limited knowledge and is aware only of some subset $B_{Ag} \subseteq \mathcal{A}$ of arguments which she can put forward in a dialogue. We assume that the attack relation is determined by the arguments, so that the knowledge of an agent can be identified with the set B_{Ag}, inducing an AF $(B_{Ag}, \mathcal{R} \cap (B_{Ag} \times B_{Ag}))$. In the remaining definitions, we assume $(\mathcal{A}, \mathcal{R})$ to be given.

Next, we need to fill in the procedural layer. The main object with which we are concerned here is a *dialogue trace*, which represents a dialogue between $\mathcal{P}$ and $\mathcal{O}$ by a sequence of moves (i. e. sets of arguments $M \subseteq \mathcal{A}$) made by $\mathcal{P}$ and $\mathcal{O}$ alternately, with $\mathcal{P}$ making the first move. Formally:

Definition 4 (Dialogue trace) *A* dialogue trace *is a finite sequence* $\pi = (M_1, \ldots, M_n)$ *s.t.* $M_1, \ldots, M_n \subseteq \mathcal{A}$. *Every* M_i *is called a* move. *We define* $A_\pi = M_1 \cup \ldots \cup M_n$ *and* $n_\pi = |M_1| + ... + |M_n|$. $\pi[n]$ *denotes the dialogue trace consisting of the first* n *moves of* π*;* $\pi[0]$ *is the empty sequence. The set of all possible dialogue traces is denoted by* S.

The rules of the dialogue are captured by the *legal move function* $\mathsf{legalmoves} : S \to 2^{2^{\mathcal{A}}}$ which returns valid follow-up moves for a particular dialogue trace. The heuristic, or strategic component is captured by an *agent model*, one for $\mathcal{P}$ and one for $\mathcal{O}$.

Definition 5 (Agent model) *An* abstract agent model Δ *is a triple* $\Delta = (\mathcal{K}, \mathsf{move}, \mathsf{upd})$ *where* $\mathcal{K}$ *is the* belief state*;* move *is a function mapping a dialogue trace and belief state to a set of moves, called the* move function*; and* upd *is a function mapping a belief state and a move to a new belief state, called the* update function.

A belief state $\mathcal{K}$ captures the agent's knowledge, utility function and opponent model. The function move returns the set of moves for the agent, given her belief state and current dialogue trace and implements the agent's strategy. We assume that an agent's move function returns only legal moves. Note that an agent may be indifferent as to which move is best and can also decide to end the dialogue, i. e., move may return multiple or zero moves. Finally, the function upd takes a belief state $\mathcal{K}$ and the move made by the opponent and yields a new belief state $\mathcal{K}'$, the idea being that moves made by the opponent may change the agent's knowledge and beliefs.

Definition 6 *A dialogue state is a pair* $(\Delta_\mathcal{P}, \Delta_\mathcal{O})$ *where* $\Delta_\mathcal{P}, \Delta_\mathcal{O}$ *are a proponent's and opponent's agent model. A dialogue trace* $\pi = (M_1, \ldots, M_n)$ *is called* valid *wrt. a legal move function* legalmoves *and a dialogue state* $(\Delta_\mathcal{P}, \Delta_\mathcal{O})$ *if and only if there exists a sequence of dialogue states* $((\Delta^0_\mathcal{P}, \Delta^0_\mathcal{O}), \ldots, (\Delta^n_\mathcal{P}, \Delta^n_\mathcal{O}))$ *with* $\Delta^i_{Ag} = (\mathcal{K}^i_{Ag}, \mathsf{upd}_{Ag}, \mathsf{move}_{Ag})$ *such that* $\Delta^0_\mathcal{P} = \Delta_\mathcal{P}$, $\Delta^0_\mathcal{O} = \Delta_\mathcal{O}$ *and, for* $i = 1, \ldots, n$*:*

1. $M_i \in \mathsf{move}_\mathcal{P}(\pi[i-1], \mathcal{K}^{i-1}_\mathcal{P})$ *if* i *is odd,*
2. $M_i \in \mathsf{move}_\mathcal{O}(\pi[i-1], \mathcal{K}^{i-1}_\mathcal{O})$ *if* i *is even,*
3. $\mathcal{K}^i_{Ag} = \mathsf{upd}(\mathcal{K}^{i-1}_{Ag}, M_i)$ *for* $Ag \in \{\mathcal{P}, \mathcal{O}\}$.

No moves can be added to a dialogue trace if an agent decides to end the dialogue. A dialogue trace is then *complete*:

Definition 7 *Let* $\pi = (M_1, \ldots, M_n)$ *be a valid dialogue trace with respect to a legal move function* legalmoves *and a dialogue state* $(\Delta_\mathcal{P}, \Delta_\mathcal{O})$. *We say* π *is* complete *if and only if there is no dialogue trace* $\pi' = (M_1, \ldots, M_n, M_{n+1})$ *which is valid with respect to* legalmoves *and* $(\Delta_\mathcal{P}, \Delta_\mathcal{O})$.

Note that, because the move function may return more than one move, there may be more than one valid and complete dialogue trace for a given pair of initial agent models. Our discourse model is thus nondeterministic with respect to how a dialogue evolves.

In the following sections we present concrete instantiations of agent models. We demonstrate these models by fixing $(\mathcal{A}, \mathcal{R})$, specifying legalmoves, and showing the resulting valid dialogue traces.

4 Agent models for strategic argumentation

In this section we present three concrete instantiations of agent models. We focus here on the belief state $\mathcal{K}$ and move function. Each model extends the expressivity of the former. We show, by example, that these extensions are necessary to properly model strategically important beliefs of an agent. In each of the three agent models, the move function is based on a variant of the M* search algorithm [Carmel and Markovitch, 1996]. We postpone the treatment of the third component of an agent model ($\mathcal{K}$, move, upd), namely the update function upd, returning to it in Section 5.

4.1 The simple agent model

The simple belief state of an agent (denoted by $\mathcal{K}_s$) consists, first, of a set $B \subseteq \mathcal{A}$ containing the arguments that the agent is aware of. The goals of the agent are encoded by the utility function u, that returns the utility (a real number) that the agent assigns to a particular dialogue trace $\pi \in S$ (cf. [Thimm and Garcia, 2010]). The agent's beliefs about the knowledge and goals of her opponent and about the beliefs about her opponent's belief about herself, etc., are modeled by simply nesting this structure, so that the third component of the simple belief state is again a simple belief state.

Definition 8 *A* simple belief state $\mathcal{K}_s$ *is a tuple* (B, u, E) *where:*

- $B \subseteq \mathcal{A}$ *is the set of arguments the agent is aware of,*
- $u : S \to \mathbb{R}$ *is the utility function,*
- $E = (B', u', E')$ *is a simple belief state called the* opponent state*, such that* $B' \subseteq B$.

The intuition behind $B' \subseteq B$ in the above stems from the common sense notion that an agent cannot have beliefs about whether or not her opponent is aware of an argument that she herself is not aware of. In other words, If an agent believes that her opponent knows argument a, then surely the agent herself also knows a. We refer to this requirement as the *awareness restriction*.

Except for this restriction, this model is the same as the one presented by Oren and Norman [Oren and Norman, 2009]. They also present a variation of the maxmin algorithm (which is in turn a variation of the M* algorithm [Carmel and Markovitch, 1996]) that determines, given a belief state $\mathcal{K}_s$ and legalmoves function, the moves that yield the best expected outcome. We can use the same approach to define our move$_s$ function. The algorithm that defines the move$_s$ function is shown in Algorithm 1. Note that the actual algorithm only needs to return the set of best moves. To simplify the algorithm, however, we define it to return both the set of best moves and the expected utility of these moves.

The algorithm works as follows. Initially the *bestMoves* is empty and *maxEU*, acting as a lower bound on the expected utility to be improved upon, is set to the utility of the current trace. For every legal move M the set of best responses of the opponent is determined on line 5. Then, in lines 6 and 7 the expected utility of M is determined by again calling the move function. On line 8 we divide the expected utility of M by the number of opponent moves, taking into account all possible equally good moves the opponent can make without double counting. Lines 9–13 keep track of the moves, considered so far, that yield maximum expected utility. Note that, if *bestMoves* is empty at the end, not moving yields highest expected utility.

Algorithm 1 $\texttt{move}_s(\pi, (B, u, E))$

1: maxEU=$u(\pi)$
2: bestMoves=$\emptyset$
3: **for all** $M \in$ legalmoves(π) **do**
4: $\quad eu = 0$
5: $\quad (oUtil, oMoves) = \texttt{move}_s((\pi, M), E)$
6: $\quad$ **for all** $M' \in oMoves$ **do**
7: $\quad\quad (nUtil, nMoves) = \texttt{move}_s((\pi, M, M'), (B, u, E))$
8: $\quad\quad eu = eu + nUtil * \frac{1}{|oMoves|}$
9: $\quad$ **if** $eu > maxEU$ **then**
10: $\quad\quad bestMoves = \emptyset$
11: $\quad$ **if** $eu \geq maxEU$ **then**
12: $\quad\quad bestMoves = bestMoves \cup \{M\}$
13: $\quad\quad maxEU = eu$
14: **return** $(maxEU, bestMoves)$

Note that we assume that the nesting of the belief state is sufficiently deep to run the algorithm. Alternatively, the algorithm can easily be adapted (as shown in [Oren and Norman, 2009]) to deal with belief states of insufficient depth, or to terminate at some fixed search depth.

In the rest of this text, we use a utility function to the effect that the dialogue is about grounded acceptance of an argument $x \in \mathcal{A}$. This has been called a *grounded game* in the literature [Modgil and Caminada, 2009]. We subtract arbitrary small values ϵ, for each move in the dialogue, capturing the idea that shorter traces are preferred, effectively driving the agents to put forward only relevant moves. Formally:

Definition 9 (Grounded game utility function) *Let* $x \in \mathcal{A}$ *and* $Ag \in \{\mathcal{P}, \mathcal{O}\}$*. The* grounded game utility function*, denoted by* $u_g^{(x,Ag)}$*, evaluated over a dialogue trace* π*, is defined by:*

$$u_g^{(x,Ag)}(\pi) = \begin{cases} v - n_\pi \epsilon & \text{if } x \text{ grounded in } F \\ -v - n_\pi \epsilon & \text{if } x \text{ attacked by gr. extension of } F \\ 0 & \text{otherwise} \end{cases}$$

where $F = (A_\pi, \mathcal{R} \cap (A_\pi \times A_\pi))$*, and* $v = 1$*, if* $Ag = \mathcal{P}$ *and* $v = -1$*, if* $Ag = \mathcal{O}$*.*

Following [Modgil and Caminada, 2009], we may refer to x as *in* if it is within the (uniquely determined) grounded extension, *out* if it is attacked by an element from this extension, and *undecided* otherwise.

For simplicity, we assume that utility functions are fixed in all models, i. e., both agents have correct beliefs about the opponent's utility function. Furthermore, the legalmoves function that we use simply forces moves to consist of a single argument, and is defined by legalmoves$_1(\pi) = \{\{x\} \mid x \in \mathcal{A}\}$.

Example 1 *Let* $(\mathcal{A}, \mathcal{R})$*, and* E, F, G *be the AF and belief states as shown in Figure 1. That is,* $E = (B, u_g^{(a,P)}, F), F = (B', -u_g^{(a,O)}, G), G = (B', u_g^{(a,P)}, F)$ *with* $B = \{a, b, c, d, e\}$ *and* $B' = \{a, b, c, d\}$*. We define the agent models* $\Gamma = (\Delta_\mathcal{P}, \Delta_\mathcal{O})$ *by* $\Delta_\mathcal{P} = (E, \text{move}_s, \text{upd})$ *and* $\Delta_\mathcal{O} = (F, \text{move}_s, \text{upd})$*, where* upd *is defined by* upd$(\mathcal{K}, M) = \mathcal{K}$*. In words,* $\mathcal{P}$ *is aware of all arguments and (correctly) believes that* $\mathcal{O}$ *is aware of only* a, b, c *and* d*. There is a single valid dialogue trace w.r.t.* Γ*, namely*

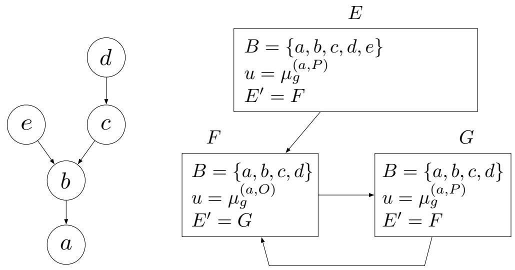

Figure 1: An AF and belief state structure.

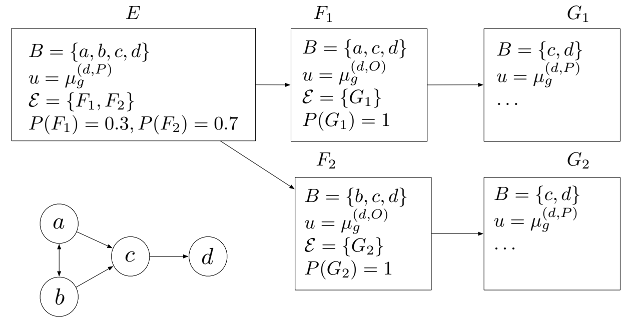

Figure 2: An AF (left) and belief state (right).

$(\{a\}, \{b\}, \{e\})$. *The first move is obvious, i. e., if $\mathcal{P}$ would not put forward the argument a under consideration, she loses because the opponent can end the dialogue. The second move b is made because $\mathcal{O}$ (not being aware of e) believes she can counter the only possible countermove after b (i. e., c) by d, resulting in a win. As it turns out, $\mathcal{P}$'s quickest way to win is by moving e and not c. This ends the dialogue, as $\mathcal{O}$ cannot increase utility by putting forward any remaining argument. The utilities of the trace are $1 - 3\epsilon$ for $\mathcal{P}$ and $-1 - 3\epsilon$ for $\mathcal{O}$.*

4.2 The uncertain agent model

A limitation of the simple agent model is that it assumes certainty about the opponent model. In the *uncertain agent model*, denoted $\mathcal{K}_u$, we capture uncertainty by assigning (non-zero) probabilities to possible opponent modes:

Definition 10 *An* uncertain belief state *$\mathcal{K}_u$ is a tuple $(B, u, \mathcal{E}, P)$ where:*

- *$B \subseteq \mathcal{A}$ is the set of arguments,*
- *$u : S \to \mathbb{R}$ is the utility function,*
- *$\mathcal{E}$ is a set of uncertain belief states (opponent belief states) such that $\forall (B', u', \mathcal{E}', P') \in \mathcal{E}$: $B' \subseteq B$,*
- *$P : \mathcal{E} \to (0, 1]$ is a probability function s.t. $\sum_{E \in \mathcal{E}} E = 1$.*

The corresponding move_u function, defined by algorithm 2, is a straightforward generalization of move_s, taking into account probabilities of possible opponent models.

Algorithm 2 $\mathtt{move}_u(\pi, (B, u, \mathcal{E}, P))$

$maxEU = u(\pi)$
$bestMoves = \emptyset$
for all $M \in \mathsf{legalmoves}(\pi)$ **do**
  $eu = 0$
  for all $E \in \mathcal{E}$ **do**
    $(oUtil, oMoves) = (\mathtt{move}_u((\pi, M), E))$
    for all $M' \in oMoves$ **do**
      $(nUtil, nMoves) = \mathtt{move}_u((\pi, M, M'), (B, u, \mathcal{E}, P))$
      $eu = eu + nUtil * P(E') * \frac{1}{|oMoves|}$
  if $eu > maxEU$ **then**
    $bestMoves = \emptyset$
  if $eu \geq maxEU$ **then**
    $bestMoves = bestMoves \cup \{M\}$
    $maxEU = eu$
return $(maxEU, bestMoves)$

Example 2 *Let $(\mathcal{A}, \mathcal{R}), E, F_1, F_2, G_1$ and G_2 be the AF and belief states as shown in Figure 2. We define the agent models $\Delta_\mathcal{P}$ by $\Delta_\mathcal{P} = (E, \mathsf{move}_u, \mathsf{upd})$ and $\Delta_\mathcal{O} = (F_1, \mathsf{move}_u, \mathsf{upd})$, where upd is defined by $\mathsf{upd}(\mathcal{K}, M) = \mathcal{K}$. In words, $\mathcal{P}$ is aware of all arguments and correctly believes $\mathcal{O}$ to be aware of c, d, but is uncertain about whether $\mathcal{O}$ knows a ($p = 0.3$) or b ($p = 0.7$). $\mathcal{O}$ in fact is aware of a and not of b. There is a single valid dialogue trace w.r.t. Γ, namely $(\{d\}, \{c\}, \{b\}, \{a\})$. Again, $\mathcal{P}$ first moves the argument d under consideration. $\mathcal{O}$, believing that $\mathcal{P}$ does not know a, replies with c, believing this will be successful. Now, $\mathcal{P}$ has to choose between putting forward a or b. Putting forward a while $\mathcal{O}$ knows b (which $\mathcal{O}$ believes to be more likely) will make d undecided. Thus $\mathcal{O}$ puts forward b. It turns out that $\mathcal{O}$ was in fact aware of a, and thus puts this argument forward, changing the status of d from out to undecided. The result is a 'tie-break', i. e., both $\mathcal{P}$ and $\mathcal{O}$ assign a utility of 0 +/- 4ϵ to the trace in which d is neither accepted nor rejected.*

4.3 The extended agent model

The two models presented so far assume that an agent cannot have beliefs about whether or not her opponent is aware of an argument that she herself is not aware of. While this is a natural assumption, it limits the kind of situations we can model. An agent can still believe that her opponent knows *some* argument, even if she is not aware of this argument herself. We model such arguments as *virtual arguments*. These are believed to exist but cannot put it forward in a dialogue. For example, if one is engaged in a dialogue with a physicist about the speed of light one may assume that the physicist has an argument for the speed of light being larger than 50 kph. However, if one is not expert in physics, the exact nature of this argument might be unknown. Furthermore, we assume that if new argument is put forward, an agent knows whether or not this argument corresponds to a virtual argument she believed to exist. That is, she can *recognize* a new argument, and map it to a virtual argument. To model this, we add a set G of virtual arguments (such that $G \cap \mathcal{A} = \emptyset$), an attack relation R between virtual arguments and regular arguments, and a recognition function rec to the belief state. Formally:

Definition 11 *Given a set of virtual arguments $\mathcal{G}$, an* extended belief state *$\mathcal{K}_e$ is a tuple $(B, u, G, R, rec, \mathcal{E}, P)$ where:*

- *$B \subseteq \mathcal{A}$ is the set of arguments,*

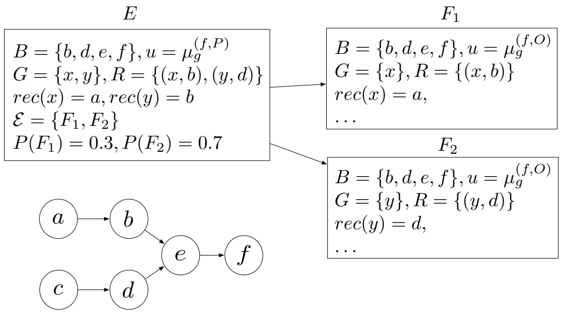

Figure 3: An AF and belief state structure.

- $u : S \to \mathbb{R}$ *is the utility function,*
- $G \subseteq \mathcal{G}$ *is the set of* virtual arguments *believed to exist,*
- $R \subseteq G \times A \cup A \times G \cup G \times G$ *is the attack relation,*
- $rec : \mathcal{A} \to 2^G$ *is the* recognition function*,*
- $\mathcal{E}$ *is a set of extended belief states, called opponent belief states, s.t.* $\forall (B', u', G', R', rec, \mathcal{E}', P') \in \mathcal{E}$*:* $B' \subseteq B$*,*
- $P : \mathcal{E} \to (0, 1]$ *is a probability function s.t.* $\sum_{E \in \mathcal{E}} = 1$.

Except for the added items, extended and uncertain belief states are similar and the $\mathtt{move}_e$ function is, except for having differently typed parameters, is identical to $\mathtt{move}_u$.

Example 3 *Let* $(\mathcal{A}, \mathcal{R}), E, F_1$ *and* F_2 *be the AF and belief states shown in Figure 3. We define the agent models* $\Delta_{\mathcal{P}}$ *by* $\Delta_{\mathcal{P}} = (E, \mathsf{move}_e, \mathsf{upd})$*, where* upd *is defined by* $\mathsf{upd}(\mathcal{K}, M) = \mathcal{K}$*. In words,* $\mathcal{P}$ *is aware of* b, d, e *and* f *and believes* $\mathcal{O}$ *is also aware of these arguments. In addition,* $\mathcal{P}$ *believes* $\mathcal{O}$ *may have counterarguments to* b *and* d*, with probability* 0.3 *and* 0.7 *respectively. In the belief state of* $\mathcal{P}$*, not being aware of* b *and* d*, they are modeled as virtual arguments, i. e.,* x *and* y *mapping to* a *and* b *respectively.*

A possible dialogue trace is $(\{f\}, \{e\}, \{b\})$*. Here,* $\mathcal{P}$ *puts forward* f *and* $\mathcal{O}$ *counters with* e*, believing this may be successful. For* $\mathcal{P}$*, the choice of whether to put forward* b *or* d *depends on her beliefs about (virtual) counterarguments.* $\mathcal{P}$ *believes it more likely (with* $p = 0.7$*) that* $\mathcal{O}$ *can counter* d *and therefore* b *is* $\mathcal{P}$*'s best move.*

5 Updating Opponent Models

When an opponent puts forward a move, an opponent model needs to be updated to take into account the knowledge conveyed by this move. We propose a number of upd functions to model such updates. (Again, we assume that utility functions are fixed and therefore do not change.)

Definition 12 *Let* $\mathcal{K}_s = (B, u, E)$ *be a simple belief state and* $M \subseteq \mathcal{A}$ *a move. The* simple update function upd_s *is defined by* $\mathsf{upd}_s(\mathcal{K}_s, M) = (B \cup M, u, \mathsf{upd}_s(E, M))$.

Note that, if $\mathcal{K}_s$ satisfies the awareness restriction then $\mathsf{upd}_s(\mathcal{K}_s, M)$ does, too.

Definition 13 *Let* $\mathcal{K}_u = (B, u, \mathcal{E}, P)$ *be an uncertain belief state and* $M \subseteq \mathcal{A}$ *a move. The* uncertain update function upd_u *is defined by* $\mathsf{upd}_u(\mathcal{K}_u, M) = (B \cup M, u, \mathcal{E}', P')$ *where*

$$\mathcal{E}' = \bigcup_{E \in \mathcal{E}} \mathsf{upd}_u(E, M) \tag{1}$$

$$P'(E) = \sum_{E' \in \mathcal{E}, \mathsf{upd}_u(e, M) = E} P(E') \qquad \textit{for } E \in \mathcal{E}' \tag{2}$$

In the above definition o we assume that arguments are observed independently of one another and, thus, probabilities stay robust in the light of observing unexpected moves. However, consider what occurs if the proponent believes that the opponent is aware of only of b with $p = 0.3$ and only of c with $p = 0.7$ how should those probabilities be adjusted when the opponent moves with argument a?

The problem is that this observation is inconsistent with the two opponent models considered possible. two ways exist to deal with this. First, we could switch to a uniform distribution, e. g. giving both states a probability of 0.5. Second, we could assume that the observations of arguments are probabilistically independent events. Taking the latter approach, we then add the observed move to every opponent model considered possible before, cf. Equation (1). Furthermore, some opponent models may collapse into one, so that we have to sum up probabilities for such states, cf. Equation (2).

Definition 14 *Let* $\mathcal{K}_e = (B, u, G, R, rec, \mathcal{E}, P)$ *be an extended belief state and* $M \subseteq \mathcal{A}$ *a move. The* extended update function upd_e *is defined via* $\mathsf{upd}_e(\mathcal{K}_e, M) = (B \cup M, u, G', R', rec, \mathcal{E}', P')$ *where*

$$G' = G \setminus \bigcup_{a \in M} rec(a) \tag{3}$$

$$R' = R \cap (G' \times B \cup B \times G' \cup G' \times G') \tag{4}$$

$$\mathcal{E}' = \bigcup_{E \in \mathcal{E}} \mathsf{upd}_3(E, M) \tag{5}$$

$$P'(E) = \sum_{E' \in \mathcal{E}, \mathsf{upd}_3(E', M) = E} P(E') \qquad \textit{for all } E \in \mathcal{E}' \tag{6}$$

The following proposition establishes a strict hierarchy of our three models w.r.t. expressivity. For example, our approaches for strategic argument selection and update coincide when restricting to less expressive models. For that, we say that a simple belief state $E = (A, u, \hat{E})$ and an uncertain belief state $E' = (A', u', \mathcal{E}, P)$ are *equivalent*, denoted $E \sim E'$, if $A = A'$, $u = u'$, $\mathcal{E} = \{\hat{E}'\}$, $P(\hat{E}') = 1$, and $\hat{E}' \sim \hat{E}$ recursively. In other words, $E \sim E'$ if E' does not provide any information beyond E. Similarly, we define equivalence to an extended belief state E'' if E'' adds no virtual arguments.

Proposition 1 *1. If* $E = (A, u, E)$ *is a simple belief state then*

(a) $E' = (A, u, \mathcal{E}, P)$ *with* $\mathcal{E} = \{E\}$ *and* $P(E) = 1$ *is an uncertain belief state and*

$$\mathtt{move}_s(\pi, E) = \mathtt{move}_u(\pi, E') \qquad \textit{for every } \pi$$

$$\mathsf{upd}_s(E, M) \sim \mathsf{upd}_u(E', M) \qquad \textit{for every } M$$

(b) $E' = (A, u, G, R, rec, \mathcal{E}, P)$ *with* $G = R = \emptyset$*,* $rec(a) = \emptyset$ *for every* $a \in A$*,* $\mathcal{E} = \{E\}$ *and*

$P(E) = 1$ *is an extended belief state and*

$$\texttt{move}_s(\pi, E) = \texttt{move}_e(\pi, E') \qquad \textit{for every } \pi$$
$$\texttt{upd}_s(E, M) \sim \texttt{upd}_e(E', M) \qquad \textit{for every } M$$

2. *If* $E = (A, u, \mathcal{E}, P)$ *is an uncertain belief state then* $E' = (A, u, G, R, rec, \mathcal{E}, P)$ *with* $G = R = \emptyset$ *and* $rec(a) = \emptyset$ *for every* $a \in A$ *is an extended belief state and*

$$\texttt{move}_u(\pi, E) = \texttt{move}_e(\pi, E') \qquad \textit{for every } \pi$$
$$\texttt{upd}_u(E, M) \sim \texttt{upd}_e(E', M) \qquad \textit{for every } M$$

Proofs are omitted due to space restrictions.

6 Implementation and Evaluation

We implemented the three different opponent models using Java in the *Tweety library for artificial intelligence*[1]. Our AF allows for the automatic generation of random abstract argumentation theories and simulates a dialogue between multiple agents. We used this AF to conduct experiments with our models and to evaluate their effectiveness in practice.

For evaluating performance we generated a random abstract argumentation theory with 10 arguments, ensuring that the argument under consideration is in its grounded extension, i. e. under perfect information the proponent should win the dialogue. However, from these 10 arguments only 50 % are known by the proponent but 90 % by the opponent. We used a proponent without opponent model and generated an extended belief state for the opponent (with maximum recursion depth 3). From this extended belief state we derived an uncertain belief state by simply removing the virtual arguments. From this uncertain belief state we derived a simple belief state by sampling a nested opponent model from the probability function in the uncertain belief state. For each belief state we simulated a dialogue against the same opponent and counted the number of wins. We repeated the experiment 5000 times, with Figure 4 showing our results. As seen, increasing the complexity of the belief state yields better overall performance. In particular, note that the difference between the performances of the simple and uncertain belief states is larger than between uncertain and extended belief states. However, this observation is highly depended on the actual number of virtual arguments used (which was around 30 % of all arguments in this experiment) and is different for larger values (due to space restrictions we do not report on the results of those experiments).

7 Related Work

Recently, interest has arisen in combining probability with argumentation. [Hunter, 2012] describes two systems which concern themselves with the likelihood that an agent knows a specific argument, and we can view the possible argument AFs that can be induced from these likelihoods as possible models of agent knowledge. [Thimm, 2012] investigates probabilistic interpretations of abstract argumentation and relationships to approaches for probabilistic reasoning. Furthermore, [Oren *et al.*, 2012] investigated strategies in such a probabilistic setting but concerned themselves with monologues rather than dialogues.

[1] http://tinyurl.com/tweety-opp

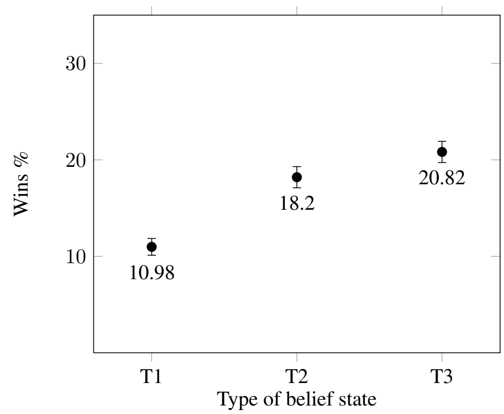

Figure 4: Performance of the simple (T1), uncertain (T2), and extended (T3) belief states in comparison (with Binomial proportion confidence intervals)

Our work concerns itself with identifying the arguments an agent should advance at any point in a dialogue. Other work in this vein includes [Oren *et al.*, 2006], which aims to minimise the cost of moves, with no concern to the opponent's knowledge, and without looking more than one step ahead when reasoning. Such a strategy can easily be encoded by our approach. By assigning probabilities to arguments, [Roth *et al.*, 2007] constructed a game tree allowing dialogue participants to maximise the likelihood of some argument being accepted or rejected. The probabilities in that system arose from *a priori* knowledge, and no consideration was given to the possibility of an opponent model.

[Rahwan and Larson, 2008; Rahwan *et al.*, 2009] consider a very different aspect of strategy, attempting to identify situations which are *strategy-proof*, that is, when full revelation of arguments is the best course of action to follow. Similarly, [Thimm and Garcia, 2010] extends that work to structured AFs and also proposes some simple dominant strategies for other specific situations. This can be contrasted with our work, where e. g. withholding information can result in a better outcome for the agent than revealing all its arguments.

8 Conclusions and Future Work

We proposed three structures for modeling an opponents belief in strategic argumentation. Our simple model uses a recursive structure to hold the beliefs an agent has on the other agent's beliefs. We extended this model to incorporate quantitative uncertainty on the actual opponent model and qualitative uncertainty on the set of believed arguments. All our models have been implemented and we tested their performance in a series of experiments. As expected, increasing the complexity of the opponent modelling structure resulted in improved outcomes for the agent.

We consider several avenues of future work. First, agents using our strategies attempt to maximise their outcome, with no consideration for *risk*. We seek to extend our work to cater for this notion by introducing second order probabilities into our system. We also intend to investigate whether virtual arguments are equivalent to a simpler system wherein no attacks between virtual arguments can exist. Furthermore, while it is difficult to obtain large scale argument graphs obtained from real world domains, we hope to validate our approach over such corpora. Finally, while our results (for clarity of presentation) focus on abstract argument, [Hadjinikolis *et al.*, 2012] has highlighted the need for strategies when structured argumentation is used. Since the work presented here can easily be extended to this domain, we are in the process of adapting our algorithms to deal with dialogues built on top of structured argumentation.

References

[Carmel and Markovitch, 1996] David Carmel and Shaul Markovitch. Learning and using opponent models in adversary search. Technical Report CIS9609, Technion, 1996.

[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):321–358, 1995.

[Grabmair and Ashley, 2010] Matthias Grabmair and Kevin D. Ashley. Argumentation with value judgments - an example of hypothetical reasoning. In Radboud Winkels, editor, *JURIX*, volume 223 of *Frontiers in Artificial Intelligence and Applications*, pages 67–76. IOS Press, 2010.

[Hadjinikolis *et al.*, 2012] Christos Hadjinikolis, Sanjay Modgil, Elizabeth Black, Peter McBurney, and Michael Luck. Investigating strategic considerations in persuasion dialogue games. In *STAIRS 2012 - Proceedings of the Sixth Starting AI Researchers' Symposium*, pages 137–148, 2012.

[Hunter, 2012] Anthony Hunter. Some foundations for probabilistic abstract argumentation. In Verheij et al. [2012], pages 117–128.

[Modgil and Caminada, 2009] Sanjay Modgil and Martin Caminada. Proof Theories and Algorithms for Abstract Argumentation Frameworks. In Guillermo Simari and Iyad Rahwan, editors, *Argumentation in Artificial Intelligence*, chapter 6, pages 105–129. Springer US, Boston, MA, 2009.

[Oren and Norman, 2009] Nir Oren and Timothy J. Norman. Arguing using opponent models. In *Proceedings of the 6th international conference on Argumentation in Multi-Agent Systems*, ArgMAS'09, pages 160–174, Berlin, Heidelberg, 2009. Springer-Verlag.

[Oren *et al.*, 2006] Nir Oren, Timothy J. Norman, and Alun Preece. Arguing with confidential information. In *Proceedings of the 18th European Conference on Artificial Intelligence*, pages 280–284, Riva del Garda, Italy, August 2006.

[Oren *et al.*, 2012] Nir Oren, Katie Atkinson, and Hengfei Li. Group persuasion through uncertain audience modelling. In Verheij et al. [2012], pages 350–357.

[Prakken and Sartor, 2002] Henry Prakken and Giovanni Sartor. *Computational Logic: Logic Programming and Beyond. Essays In Honour of Robert A. Kowalski, Part II*, volume 2048 of *LNCS*, pages 342–380. Springer-Verlag, 2002.

[Rahwan and Larson, 2008] Iyad Rahwan and Kate Larson. Mechanism design for abstract argumentation. In *Proceedings of AAMAS 2008*, 2008.

[Rahwan *et al.*, 2009] Iyad Rahwan, Kate Larson, and Fernando A. Tohmé. A characterisation of strategy-proofness for grounded argumentation semantics. In *IJCAI 2009, Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 251–256, July 2009.

[Roth *et al.*, 2007] B. Roth, R. Riveret, A. Rotolo, and G. Governatori. Strategic Argumentation: A Game Theoretical Investigation. In *ICAIL '07: Proceedings of the 11th International Conference on Artificial Intelligence and Law*, pages 81–90, New York, NY, USA, 2007. ACM Press.

[Thimm and Garcia, 2010] Matthias Thimm and Alejandro J. Garcia. Classification and Strategical Issues of Argumentation Games on Structured Argumentation Frameworks. In Wiebe van der Hoek, Gal A. Kaminka, Yves Lespérance, Michael Luck, and Sandip Sen, editors, *Proceedings of the Ninth International Joint Conference on Autonomous Agents and Multi-Agent Systems 2010 (AAMAS'10)*, May 2010.

[Thimm, 2012] Matthias Thimm. A probabilistic semantics for abstract argumentation. In *Proceedings of the 20th European Conference on Artificial Intelligence (ECAI'12)*, August 2012.

[Toniolo *et al.*, 2011] Alice Toniolo, Timothy J. Norman, and Katia P. Sycara. Argumentation schemes for collaborative planning. In David Kinny, Jane Yung jen Hsu, Guido Governatori, and Aditya K. Ghose, editors, *PRIMA*, volume 7047 of *Lecture Notes in Computer Science*, pages 323–335. Springer, 2011.

[Verheij *et al.*, 2012] Bart Verheij, Stefan Szeider, and Stefan Woltran, editors. *Computational Models of Argument - Proceedings of COMMA 2012, Vienna, Austria, September 10-12, 2012*, volume 245 of *Frontiers in Artificial Intelligence and Applications*. IOS Press, 2012.