

Empirical Evaluation of Strategies for Multiparty Argumentative Debates

Dionysios Kontarinis¹, Elise Bonzon¹, Nicolas Maudet², and Pavlos Moraitis¹

¹ LIPADE, Université Paris Descartes,

dionysios.kontarinis@parisdescartes.fr,
{elise.bonzon,pavlos}@mi.parisdescartes.fr

² LIP6, Université Pierre et Marie Curie, nicolas.maudet@lip6.fr

Abstract. Debating agents have often different areas of expertise and conflicting opinions on the subjects under discussion. They are faced with the problem of deciding how to contribute to the current state of the debate in order to satisfy their personal goals. We focus on target sets, that specify minimal changes on the current state of the debate allowing agents to satisfy their goals, where changes are the addition and/or deletion of attacks among arguments. In this paper, we experimentally test a number of strategies based on target sets, and we evaluate them with respect to different criteria, as the length of the debate, the happiness of the agents, and the rationality of the result.

1 Introduction

In recent years, the study of the collective aspects of argumentation (which can now be increasingly experienced on-line [1]), has seen a surge of interest in AI. Such settings raise new challenges for argumentation theory [2]. The object constructed by a group of agents is a *weighted argumentation system* [3], where a natural interpretation of the weights attached to an edge is that it reflects the number of agents who have committed to a given attack, or the aggregated expertise of those agents [4], as we shall also assume here. New semantics have been proposed to account for the social nature of argumentation and its specific use in a context where votes can be cast on top of arguments (and relations among them), either sticking to the framework initially set up by Dung [5], see *e.g.* [6], or departing from it [7, 8].

Debates in online settings are incrementally built, with agents adding new arguments, attacks, and casting new votes in response to the opinion voiced by others. In practice such debates may be (more or less flexibly) regulated, to ensure that they remain focused, and that some fairness is guaranteed among the different agents. One thing that is missing though is a study of the dynamics of debates regulated by such protocols: it is not clear how strategies used by agents would change the outcome of debates. In [9] a very simple dynamic is investigated, based on a direct notion of relevance inspired by [10], and it is shown that in the absence of coordination and with a myopic behavior, agents can actually play against their own interest, leading to undesirable results. This justifies the fact that some “guidance” might be useful to agents, without assuming though any sort of explicit coordination among agents. Recently, the notion

of *target sets* has been proposed in the literature [11, 12]. Roughly speaking, a target set specifies the minimal (sets of) moves which would achieve the argumentative goal of a given side of the debate, *provided the debate remains in its current state*. The intuition is that agents should be better off focusing their moves on target sets. One challenge though is that target sets may prescribe more than one move for agents to play, and that it is impossible to assume that agents will have the opportunity to completely “control” a target set.

In this paper we experimentally investigate how well strategies based on target sets behave. We study a number of dynamics, of increasing complexity, where the notion of target set is thoroughly exploited. Our experimental results show in particular that the use of these sophisticated strategies provides an advantage to the side using it, and that it shortens the length of debates.

The rest of this paper is as follows. Section 2 provides the necessary background, introducing the different elements composing the “gameboard” of the debate. Section 3 recalls the definition of target sets. Section 4 presents a protocol, and Section 5 presents a study of different strategies of increasing complexity, based on this notion of target sets. These strategies are experimentally compared in Section 6. Section 7 concludes.

2 Argumentative Debates Featuring Conflicting Expert Opinions

The aim of this work is to study argumentative debates among expert agents. We consider an arbitrary number of participating agents, each of them having a private argumentation system. For the sake of simplicity we assume that all agents have the same set of arguments, but they can disagree on the validity of the attacks between those arguments. Each argument concerns a finite set of topics, and the agents are experts on a subset of these topics. The debate is about the status (wrt a given semantics) of a single argument, called *issue*. The agents vote on the attacks involved in the computation of the status of the issue, on a specific common system called Gameboard. The objective of each agent is to have the status of the issue in his private argumentation system be the same as the status of the issue on the Gameboard, at the end of the debate.

2.1 Modelling the Participants

A finite set of agents, denoted Ag , take part in a debate. Each agent $i \in Ag$ has a private Dung argumentation system [5], where the exact structure of the arguments is unspecified. All agents share the same set of arguments A , but they may disagree on the attacks between them. For this reason we introduce the notion of **master argumentation system** which contains all attacks on which the agents agree, as well as all attacks on which they disagree. The attacks on which the agents agree are called fixed (or undeniable). Private argumentation systems of agents inherit fixed attacks appearing in the master AS. More formally:

Definition 1. An **argumentation system (AS)** is a pair $\langle A, R \rangle$ of a set A of arguments and a binary relation R on A called the **attack relation**. $\forall a, b \in A, aRb$ (or $(a, b) \in R$) means that a **attacks** b .

Given a **master argumentation system** $AS = \langle A, R \rangle$ and $R^* \subseteq R$ a set of **fixed attacks**, an agent $i \in Ag$ is equipped with a private argumentation system denoted $AS_i = \langle A, R_i \rangle$, such that $R_i \subseteq R$ and $R^* \subseteq R_i$. Attacks in $R \setminus R^*$ are called **debated attacks**.

In Dung's framework, the *acceptability* of an argument depends on its membership to some sets, called extensions. These extensions characterize collective acceptability. Several *semantics for acceptability* have been defined in [5]. In what follows, we concentrate on the notion of *grounded semantics*, which can be defined as follows:

Definition 2. Let $AS = \langle A, R \rangle$ and $C \subseteq A$. The set C is **conflict-free** iff $\nexists a, b \in C$ such that aRb . C **defends** an argument a iff $\forall b \in A$ such that bRa , $\exists c \in C$ such that cRb . C is a **grounded extension** of AS iff C is the least fixed point of the characteristic function of AS ($F: 2^A \rightarrow 2^A$ with $F(C) = \{a \mid C \text{ defends } a\}$).

Intuitively, a *grounded extension* contains all arguments which are not attacked, as well as the arguments which are defended (directly or not) by non-attacked arguments. There always exists a unique grounded extension which, however, might be the empty set. Thus, all the debating agents know, at every time of the debate, which arguments are accepted and which are not. We shall denote by $Gr(AS)$ the grounded extension of the system AS .

Example 1. Let a master system $AS = \langle A, R \rangle$, with $A = \{a, b, c, d\}$, $R = \{(a, b), (b, c), (d, c)\}$ and $R^* = \{(b, c)\}$. This system can be represented as follows, where fixed attacks are represented by thick arrows, and debated attacks by simple arrows.



Let three agents, such that $AS_1 = \langle A, R_1 \rangle$, with $R_1 = \{(a, b), (b, c)\}$; $AS_2 = \langle A, R_2 \rangle$, with $R_2 = \{(b, c), (d, c)\}$ and $AS_3 = \langle A, R_3 \rangle$, with $R_3 = \{(b, c)\}$. We have $Gr(AS_1) = \{a, c, d\}$, $Gr(AS_2) = \{a, b, d\}$ and $Gr(AS_3) = \{a, b, d\}$.

Each argument is associated with a set of keywords specifying which topics this argument is about. This is common practice in systems like the ones in [2, 4]. We assume that there is a fixed set of topics, denoted T , and every argument concerns a subset of T .

Definition 3. Let T be the set of topics. The set of **topics of an argument** $a \in A$ is given by function $top(a) \subseteq T$. The set of **topics of an attack** $(a, b) \in R$ is given by function $top(a, b) = top(a) \cup top(b) \subseteq T$. The **expertise of agent** $i \in Ag$ is given by $exp(i) \subseteq T$.

2.2 Modelling the Gameboard

Inspired from [9], we use a central structure called gameboard (GB in short). The gameboard stores all the opinions expressed by the agents during the debate and aggregates

them, giving rise to a single argumentation system, which will allow us to draw the debate's conclusions. An essential element in the debates we consider, is that agents may disagree on the existence of some attacks. Thus, an agent can vote either for, or against the existence of an attack. The role of the GB is to gather and aggregate all the votes cast during the debate. Moreover, the voters' relevant expertise will play a crucial role in determining the result of the aggregation. In the rest of the paper, we assume that $AS = \langle A, R \rangle$ is a master argumentation system, R^* is the set of fixed attacks, T a set of topics, and Ag is a set of agents, such that $\forall i \in Ag, AS_i = \langle A, R_i \rangle$.

Definition 4. A *vote*, also called a *move*, is a tuple $\langle (a, b), s, i \rangle$ where $(a, b) \in R \setminus R^*$ is the debated attack³ concerned by the vote, $s \in \{-1, +1\}$ is the sign of the vote, and $i \in Ag$ is the voter.

A positive vote by an agent means that he supports that the attack does hold, while a negative vote means that he supports the opposite.

Let $(a, b) \in R \setminus R^*$, then $eval(a, b)$ is the **evaluation vector of** (a, b) . This vector contains $|top(a, b)|$ elements.

Definition 5. Let $(a, b) \in R \setminus R^*$ with $top(a, b) = \{t_1, \dots, t_n\}$. The evaluation vector of (a, b) is denoted $eval(a, b) = \langle v_{t_1}, \dots, v_{t_n} \rangle$. The value $v_{t_i} \in \mathbb{Z}, \forall t_i \in top(a, b)$, depends on the voters' expertise in t_i . Whenever a vote $\langle (a, b), s, i \rangle$ is cast by agent i , then the vector $eval(a, b) = \langle v_{t_1}, \dots, v_{t_n} \rangle$ is updated into:
 $\langle v_{t_1} + s \times |\{t_1\} \cap exp(i)|, \dots, v_{t_n} + s \times |\{t_n\} \cap exp(i)| \rangle$.

Example 1, cont. Let $T = \{t_1, t_2, t_3, t_4\}$, with $top(a) = \{t_1, t_2, t_3\}$, $top(b) = \{t_2\}$, $top(c) = \{t_2, t_3\}$, $top(d) = \{t_4\}$. Also, let $exp(1) = \{t_1, t_2\}$, $exp(2) = \{t_2, t_3\}$ and $exp(3) = \{t_1, t_4\}$. Initially, no votes have been cast on any attack belonging to the master AS. Agent 1 votes for attack (a, b) . We then have, as $top(a, b) = \{t_1, t_2, t_3\}$ and $exp(1) = \{t_1, t_2\}$, that $eval(a, b) = \langle 1, 1, 0 \rangle$. Next, agent 2 votes against attack (a, b) and for attack (d, c) . We then have $eval(a, b) = \langle 1, 0, -1 \rangle$ and $eval(d, c) = \langle 1, 1, 0 \rangle$. Finally, agent 3 votes also against (a, b) . We have $eval(a, b) = \langle 0, 0, -1 \rangle$.

Given an evaluation vector $eval(a, b)$, we can decide whether attack (a, b) should be accepted or rejected. We underline that there exist various methods to obtain such a verdict, given an evaluation vector. Here we use a simple method taking into account all the elements of an evaluation vector and using a simple sum.

Definition 6. Let $(a, b) \in R \setminus R^*$ and let $eval(a, b) = \langle v_{t_1}, \dots, v_{t_n} \rangle$ be its evaluation vector. The **verdict on** (a, b) , denoted $verdict(a, b) \in \{true, false\}$ is computed as follows: $verdict(a, b) = true$ iff $\sum_{i=1}^n v_{t_i} > 0$, $verdict(a, b) = false$, otherwise.

In other words, the verdict on an attack is positive if the aggregated relevant expertise of agents having voted for the attack is strictly greater than the aggregated relevant expertise of agents having voted against it. Otherwise, the verdict on the attack is negative.


Let us now see how a gameboard is defined. Its main feature is a set containing the evaluation vectors of all the possible attacks.

³ We assume that the agents cannot vote on the attacks which are fixed in the master system.

Definition 7. A **gameboard** is a triplet $GB = \langle A, R, Eval \rangle$, where A is the set of arguments shared by all agents, R is the set of attacks and $Eval$ is the set of evaluation vectors of $R \setminus R^*$.

Let $AS_{GB} = \langle A, R_{GB} \rangle$ be the **argumentation system of the GB**, such that $R_{GB} = \{(a, b) \in R \setminus R^* \mid verdict(a, b) = true\} \cup R^*$.

Example 1, cont. Let the gameboard $GB = \langle A, R, Eval \rangle$ with the function $Eval$ defined as previously. We have $verdict(a, b) = false$ and $verdict(d, c) = true$. This gameboard can be represented as follows:



Fixed attacks are represented by thick arrows, attacks with true verdict by simple arrows, and attacks with false verdict by dashed arrows. The argumentation system $AS_{GB} = \langle A, R_{GB} \rangle$ contains thus only the fixed attacks and the debated attacks with true verdict.



2.3 Merged System

When participating to the debate, agents are assumed truthful, and they cannot vote for (resp. against) an attack if they think that it does not (resp. does) hold. Certainly, absolute truthfulness is not often encountered in real-life debates, but it is an assumption preventing the agents from stating anything that may help them in the debate. A more refined approach, left for future work, would be to define a set of beliefs (in our case attacks) upon which an agent is able to lie, if he considers it favorable at some point. This kind of situation has already been studied by Rahwan et al [13]. In this work, they introduced a formal argumentation theory, namely ArgMD, in which an agent may hide an argument or lie about arguments.

On the other hand, we allow agents to *not* express their opinion on some attacks, because that could harm their purpose, or make them disclose information they wish to hide. This is related to the notion of *dishonest arguments*, that has been studied by Caminada in [14] and by Sakama in [15]. We thus need a way to compare the results obtained in our debates with a collective view of the argumentation systems of the agents. We rely on two different notions. The first one is the notion of *merged* argumentation system [16]. In the specific case we discuss here, it turns out that a meaningful way to merge is to take the vote of all agents on all attacks in $R \setminus R^*$.

Definition 8. Let $AS = \langle A, R \rangle$ be a master AS and Ag be a set of agents. The **merged argumentation system** is $AS_{Ag}^M = \langle A, R^M \cup R^* \rangle$ where $R^M \subseteq R$ and $aR^M b$ iff $verdict(a, b) = true$ when all the agents in Ag have voted on (a, b) .

Another notion which can be useful for analyzing the collective view of the debate is the one of *happiness*: we could want to see a majority of agents satisfied at the end of the debate, in the sense that they agree with the status of a specific argument, called *issue of the debate*.

Definition 9. Let Ag be a set of agents and let, $\forall i \in Ag$, AS_i denote i 's private AS. Also, let $c \in A$ be an argument called issue of the debate. The **majority result** is denoted $majIn(c) \Leftrightarrow |\{i \in Ag \mid c \in Gr(AS_i)\}| \geq |Ag|/2$, and it is denoted $majOut(c) \Leftrightarrow |\{i \in Ag \mid c \notin Gr(AS_i)\}| > |Ag|/2$.

Note that ties for the majority are broken in favour of the agents who want to see the issue in the grounded extension.

3 Focus on Minimal Changes

At this point we turn our attention to possible strategic considerations of agents in this type of debates. What are the attacks of the gameboard on which the voting agents should focus and try to add/remove? The aim of the analysis that follows is to provide insight on how to vote in order to achieve a goal. We will focus on target sets [11, 12], which represent the minimal change on an argumentation system, achieving an argumentative goal.

A target set is a minimal set of *actions* on an argumentation system allowing to achieve a given goal. Please note beforehand that an action is *not* the same notion as a vote on the gameboard: here we assume that an action *changes* the verdict of some attacks, whereas a vote on the gameboard does not necessarily do the same (as an agent does not always have a sufficiently high expertise). Thus, we focus, on the following definition, on the verdict of the attacks in the gameboard, and not on the exact value of the eval functions.

Definition 10. Let $GB = \langle A, R, Eval \rangle$ be the gameboard at a given time. An **action on GB** is a set of atoms $m = \{((x, y), s) \mid (x, y) \in (R \setminus R^*), s \in \{+, -\}\}$, such that $\forall((x, y), +) \in m$, $verdict(x, y) = false^4$, and $\forall((x, y), -) \in m$, $verdict(x, y) = true$.


The **resulting GB** after playing an action m , is denoted $\Delta(GB, m) = \langle A, R, Eval^m \rangle$, such that $\forall(x, y) \in R \setminus R^*$:

1. $verdict^m(x, y) = true$ iff either $verdict(x, y) = true$ and $((x, y), -) \notin m$, or $((x, y), +) \in m$.
2. $verdict^m(x, y) = false$ iff either $verdict(x, y) = false$ and $((x, y), +) \notin m$, or $((x, y), -) \in m$.

Here is an example showing how an action modifies a system.

Example 1, cont. We take the same gameboard GB as defined previously. If we play the action $m = \{((a, b), +), ((d, c), -)\}$ on GB , we obtain the following system $\Delta(GB, m)$:

⁴ That is, $Eval(x, y)$ is such that $verdict(x, y) = false$



In order to define the notion of target set, we first need to provide the definition of the goal of a debate.

Definition 11. Let GB be a gameboard, $AS_{GB} = \langle A, R_{GB} \rangle$ be its system, and $d \in A$ be the issue. The goal g_d^+ (resp. g_d^-) is **satisfied** in GB iff $d \in Gr(AS_{GB})$ (resp. $d \notin Gr(AS_{GB})$).

Definition 12. Let GB be a gameboard and let g be a goal. m is a **successful action** on GB for goal g iff g is satisfied in the resulting gameboard $\Delta(GB, m)$. We denote $\mathbb{M}(GB, g)$ the set of all successful actions on GB for goal g . m is a **target set** on GB for goal g iff m is a minimal (w.r.t. \subseteq) element of $\mathbb{M}(GB, g)$. We denote $\mathbb{T}(GB, g)$ the set of all target sets on GB for goal g .

In [12] we studied the evolution of target sets when changes occur on a system (that is, when an action is done). We have shown that if an agent plays an action which does not contain any atom of any target set, then the target sets of the new gameboard will “grow”, and it will become harder (or at least not easier) to satisfy the goal under consideration. On the other hand, if an agent plays in a target set, then that target set will “shrink”, regardless of what happens to other target sets. In that sense, at least one “path” towards the satisfaction of the goal becomes shorter, while this is not the case if we do not play on any target set.

However, as we will show in Section 5, at some point during a debate, an agent may be better off playing a move outside target sets, as this may incite his opponents to play a move which will backfire. We will also propose different strategies with which agents can choose their moves, focusing on target sets, and we will experimentally test them.

4 A Debate Protocol

In this section we define a specific debate protocol. The agents focus on the status (under the grounded semantics) of a single argument $d \in A$, which is the issue of the debate. The goal of an agent $i \in Ag$ is therefore to have the issue’s status be the same, on the GB and in his private system, at the end of the debate. We can therefore distinguish two groups of agents: the agents of the group *PRO* (resp. *CON*) who have (resp. do not have) the issue in the grounded extension of their systems. An advantage of using grounded semantics is that the grounded extension is easy to calculate and it is always unique. Therefore, every agent is either *PRO* or *CON*, and at every point of the debate the issue is either accepted or rejected.

The protocol proceeds in timesteps. Let $GB^t = \langle A, R, Eval^t \rangle$ denote the gameboard at timestep t . At $t = 0$, we have $GB^0 = \langle A, R, Eval^0 \rangle$, with $\forall (a, b) \in R \setminus R^*$, $Eval^0(a, b) = \langle 0, \dots, 0 \rangle$. Recall that attacks in R^* are fixed in the system, and cannot be modified (so, they are not associated to any evaluation vector), whereas attacks in any $R' \subseteq (A \times A) \setminus R$ cannot be added. In order to ensure the termination of our protocol, we assume that an agent cannot vote on the same attack twice. To account for

this, each agent $i \in Ag$ is equipped with a set $HV_i^t \subseteq R$ which contains all the attacks agent i has voted on, until timestep t . The protocol is defined by the following:

- **Participants:** A finite set of agents Ag , each one being either *PRO* or *CON*, according to his opinion on the issue’s status.
- **Turntaking:** Round-robin. The token is given to each agent, in turn, and comes back to the first agent once all agents have played.
- **Permitted moves:** Agent i at timestep t can either:
 - Vote on $\langle (a, b), +, i \rangle$, if $(a, b) \in R_i \setminus R^*$ and $(a, b) \notin HV_i^t$
 - Vote on $\langle (a, b), -, i \rangle$, if $(a, b) \in R \setminus R_i$ and $(a, b) \notin HV_i^t$
 - Play a **pass move** (giving the token to the next agent).
- **Stopping condition:** $|Ag|$ pass moves have been played in a row.
- **Winning condition:** Once the debate has stopped, all *PRO* (resp. *CON*) agents win iff the issue belongs (resp. does not belong) to the grounded extension of the argumentation system of AS_{GB} .

5 Strategies

When having the token, an agent can vote on any of the attacks under discussion, but which one should he choose? In general, a *strategy* states, for each agent, what move should be uttered next in the course of the debate. When a strategy returns a single move, we say it is *deterministic*. Depending on the information required to take this decision, we can distinguish different kind of strategies:

- *(k)-history-based strategies:* the strategy selects moves based on the last k moves uttered in the debate, noted *h(k)-strategies*. For instance:
“If someone just attacked argument a , I will try to defend it.”
- *(k)-state-based strategies:* the strategy selects moves based on the last k states of the gameboard, noted *s(k)-strategies*. For instance:
“If $a \in Gr(AS_{GB})$, then I will utter the attack (d, a) .”

We say that a strategy s has a richer information basis than a strategy s' (noted $s \triangleright s'$) when it uses more information to select the next moves. Observe that, for a round t , both *h(t)-strategies* and *s(t)-strategies* are fully expressive, since they can capture the whole history of the debate so far. Note also that *h(k)-strategies*, based on

$$[GB^t, GB^{t-1}, GB^{t-2}, \dots, GB^{t-k}]$$

could as well be expressed as a strategy based on the single state GB^{t-k} , together with the k last moves. Also, when $t > t'$, a t -state-based strategy has a richer information basis than a t' -history-based-strategy. Finally, for the same k , state-based and history-based strategies are incomparable: for instance, a strategy based on the last state of the gameboard may capture intuitively more information than a strategy based on the last move, but it misses the information of what was the last move uttered.

In what follows, we study a natural class of *s(1)-strategies*, as we define strategies based on the computation of the target sets of the last GB . We also make the assumption


that all agents from one side (PRO or CON) use the same strategy. This facilitates the analysis, but constitutes of course a simplification. Moreover, we assume that the agents cannot disclose their private argumentation systems. Thus, agents do not have any knowledge on the other agents' private systems. As said before, the analysis of target sets and their properties leads us naturally to think that agents would profit from focusing on attacks of target sets, as it is the fastest and most economical way to achieve a goal.

5.1 Lack of Dominance and Equilibrium Guarantees

Dominance. One may wonder whether “playing within target sets” is a dominant strategy, that is, whether agents can never be better off playing a different strategy, whatever the strategy of the other party is. Note first that “playing within target sets” does not constitute a single strategy, but instead a class of strategies, in fact a subclass of $s(I)$ -strategies. So when say “a dominant strategy”, we abuse language and mean *any* strategy belonging to this class. This turns out to be a too demanding notion, because the strategy of the other player can be of any kind, in particular, it may be such that moves played outside a target set will precisely be the moves required to lead to a winning result.

This may be illustrated in the following scenario:

Example 1, cont. Assume that we are in the beginning of the debate, and no moves have been played yet. Let the gameboard be $GB^0 = \langle A, R, Eval^0 \rangle$. We have $eval(a, b) = \langle 0, 0, 0 \rangle$, so $verdict(a, b) = false$, and $eval(d, c) = \langle 0, 0, 0 \rangle$, so $verdict(d, c) = false$. Now, we add a new argument e in A (referring to topic t_2), as well as the attack (b, e) in R , with $eval(b, e) = \langle 0 \rangle$, so $verdict(b, e) = false$. This new gameboard is represented as follows:




Agent 1, who belongs to *PRO*, focuses on target set $\{((a, b), +)\}$, as adding (a, b) will make c accepted. However, it is impossible for him alone to impose the attack (a, b) , as both agents 2 and 3 will disagree on its existence.

But now suppose that an agent of the *CON* team (eg. agent 2) is very picky on the issue of the (new) argument e , and he has a strategy which says: “If e is attacked, then I will defend e ”⁵. Of course this strategy is not directly focused on the topic of the debate (which is c), but this kind of rhetorical move is common in real-life argumentation. In this case, agent 1 has an incentive to play move $((b, e), +)$ (provided that he is able to), and lure agent 2 in responding with $((a, b), +)$. This way, thus not focusing always on target sets, agent 1 can eventually make c accepted and win the debate.

⁵ For the sake of simplicity, let us assume that here agent 2 may violate his truthfulness.

Symmetric Equilibrium. The previous example showed that not focusing on target sets may in some cases lure the other group to make a “bad” move. Of course this relies on the rather artificial construction consisting of an agent playing a somewhat irrational strategy. We may then ask whether a weaker property can be guaranteed: is it the case that, *if the other agent follows a strategy consisting of playing within target sets*, then agents of the other side will not have an incentive to play differently, ie. whether this constitutes a symmetric equilibrium. The following example, shows that this is not the case either.

Example 2. Four agents have the following argumentation systems (we assume for the sake of simplicity that the arguments concern the same topic, and that all agents are expert on this topic):



The dialogue’s issue is argument c . We have $CON = \{a_1, a_2\}$, $PRO = \{a_3, a_4\}$. If both teams of agents play only in the targets sets, agents in PRO cannot win: at the beginning of the debate, agents in CON have two target sets $\{(b, c), +\}$ and $\{(a, c), +\}$. If they vote on (a, c) , agents in PRO will be able to remove that attack (by voting twice if it is necessary). The remaining target set for CON will then be $\{(b, c), +\}$. Once CON agents vote on (b, c) , agents in PRO will have two target sets: $\{(b, c), -\}$ and $\{(a, b), +\}$. Assume that a_4 votes against (b, c) . Then agents in CON can vote again to reinstate it. Agents in PRO have then one remaining target set: $\{(a, b), +\}$. Once this vote is cast, the target set for CON is $\{(a, b), -\}$. a_2 votes against (a, b) , and the agents in PRO cannot do anything else. In this case, PRO agents cannot win the debate.

Assume now that agents in PRO do not play only in the target sets. As previously, at the beginning, agents in CON have two target sets, $\{(b, c), +\}$ and $\{(a, c), +\}$. Once again, they can vote on (a, c) but these votes will be removed by agents in PRO. Once CON agents vote on (b, c) , assume that a_4 votes on (b, d) . The target set for CON is empty (as their goal is satisfied). a_4 , for the PRO team, can play once more, so he chooses to add (d, c) , and then to remove (b, c) . The group CON has now two target sets, $\{(a, b), +\}$ and $\{(b, c), +\}$. Assume that a_1 votes for (a, b) . Agents in PRO have now two target sets, $\{(a, b), -\}$ and $\{(d, c), -\}$. If a_3 votes against (d, c) , agents in CON will have one target set, $\{(a, b), -\}, \{(b, c), +\}$. Assume that a_2 votes against (a, b) , and after everybody passes, he votes again for (b, c) . a_3 can now vote for (a, b) . Agents in CON cannot do anything else, as a_2 has already voted against (a, b) once. PRO wins the debate.

All in all, playing in target sets looks intuitively like a good strategy, but it seems difficult to obtain theoretical guarantees. This leads us to study it experimentally.

5.2 Strategies Based on Target Sets

Here we define 5 strategies, from the simpler to the more complex, mainly focusing on target sets. Strategy 0 is the exception, as it is a *random strategy*, which will allow us to assert that playing in the target sets is useful. We remind that, at any timestep, an agent is winning (resp. losing) the debate if the status of a given issue is the same (resp. is not the same) both in his private system and in the argumentation system associated to the GB. Note that when there are no available moves for an agent (we remind that an agent cannot vote on the same attack twice), that agent obligatorily passes.

Strategy 0: This is a random strategy, where (1) if the agent is winning, then he plays pass. (2) otherwise, he votes randomly on an attack on the gameboard.

Strategy 1: The idea of this strategy is to allow only agents who are not satisfied by the current state of the gameboard to vote. Moreover, these agents can only vote if they can change the status of the issue (and thus, if they can change the verdict of an attack belonging to a target set of cardinality 1).⁶ More precisely: (1) if the agent is winning, then he plays pass. (2) otherwise, the agent can only vote on an attack if this vote allows to change the status of the issue.

Strategy 2: This strategy improves the previous one by allowing agents to vote on a target set of cardinality greater than 1: an agent can vote on an attack if he can change its verdict, but this vote does not have to change the status of the issue. More precisely: (1) if the agent is winning, then he plays pass. (2) otherwise, the agent can only vote on an attack if this attack belongs to a target set, and if this vote allows to change the verdict on this attack.

Strategy 3: This strategy allows an agent to vote on an attack belonging to a target set, even if he cannot change the verdict on this attack. More precisely: (1) if the agent is winning, then he plays pass. (2) otherwise, the agent can vote on any attack belonging to a target set (towards changing the verdict).

Strategy 4: This strategy improves the previous one by allowing a winning agent to play a move which renders the goal of the other team more difficult to be reached. More precisely: (1) if the agent is winning, then he can vote on an attack which belongs to a target set for the goal of the other team and “reinforce” it.⁷ (2) otherwise, the agent can vote on any attack belonging to a target set (towards changing the verdict).

As we can have several target sets, and several actions in a target set, an agent can have several possible votes for each of these strategies. We thus introduce three heuristics to help an agent to choose which vote to cast.

5.3 Heuristics

An agent can compute a set of possible votes, using any of the above strategies. Then, he can either randomly choose a vote among them, or use a more subtle heuristic. We have defined three heuristics which can be used for filtering the initial set of possible votes.

⁶ Note that this strategy is the one studied in [9].

⁷ And thus making it more difficult for the other team to change the verdict on this attack.

- **Heuristics A:** the agent randomly chooses a possible vote.
- **Heuristics B:** the agent filters out all possible votes on non-minimal (wrt. cardinality) target sets ⁸. Then, he randomly chooses a vote.
- **Heuristics C:** the agent filters out all possible votes on non-minimal (wrt. cardinality) target sets. If he can change the verdict of an attack among the remaining ones, he filters-out all the attacks he cannot change. Then, he randomly chooses a vote.

5.4 Strategy and Debate Profiles

Coupling a strategy with a heuristics gives us a specific **strategy profile**. As Strategy 0 does not use target sets, it can not be coupled with any heuristics. Also, in Strategy 1 an agent can only vote on an attack if it belongs to a target set of cardinality 1 and he can change its verdict, so it does not make any sense to associate Strategy 1 with heuristics B or C. In the same way, in Strategy 2 an agent can only vote on an attack if he can change its verdict, so it does not make sense to couple Strategy 2 with heuristics C. We thus have the following strategy profiles to consider (the number indicates the strategy type and the capital letter the heuristics): $SP = \{0, 1, 2A, 2B, 3A, 3B, 3C, 4A, 4B, 4C\}$.

We assume that the agents of the same group (PRO or CON) are using the same strategy profile during a debate. This is done in order to draw more easily conclusions on how the strategy profiles fare against each other. We can thus introduce the notion of **debate profile**. A debate profile is defined as a couple (SP_{PRO}, SP_{CON}) with $SP_{PRO}, SP_{CON} \in SP$. It indicates that all agents in the PRO (resp. CON) group are using the strategy profile SP_{PRO} (resp. SP_{CON}). Since there are 10 strategy profiles, there exist $10 \times 10 = 100$ different debate profiles. In the following, we first examine Strategy 0, and then we turn our attention to the 9 other strategy profiles which use target sets (thus on their corresponding $9 \times 9 = 81$ debate profiles).

6 Experimental Results - Discussion

We have implemented in Java the debate framework presented in the previous sections and performed a number of experiments.

6.1 Generating Debate Configurations

In order to perform an important number of debates, our program is able to generate different *debate configurations*. A configuration consists on three elements: the set of all topics, a master argumentation system AS , and a set of agents with their private systems and their expertise. In our experiments we made the following choices:

Topics: We have $|T| = 6$ topics.

⁸ For example, if an agent can vote on two attacks, the first being in a target set of cardinality 1, and the second in a target set of cardinality 2, then he will filter out the second option.

Master argumentation system: Every generated argumentation graph contains $|A| = 20$ arguments, each one randomly attached to one or two topics. The graph has a density of attacks equal to 0.1. Among the attacks, 10 are debated, and thus belong into $R \setminus R^*$.⁹ Finally, the issue is randomly chosen among the arguments in A .

Agents: Each debate involves 10 agents. Each of them is expert in one, two, or three topics randomly chosen. The AS_i of each agent includes all the attacks in R^* , whereas each debated attack in $R \setminus R^*$ belongs to R_i with a 50% probability.

6.2 The Debates

A number of configurations were randomly generated, using the above parameter values. When the difference in the number of agents in groups *PRO* and *CON* was important, the debates were trivial, as the majority easily won. The impact of the groups' size difference is now studied in more detail: we randomly generated 10 configurations for each combination of *PRO* and *CON* cardinalities (so, 10 configurations with 9 *PRO* and 1 *CON* agents (denoted 9/1), then 10 configurations with 8 *PRO* and 2 *CON* agents (denoted 8/2), and so on (7/3, 6/4, 5/5, 4/6, 3/7, 2/8, 1/9)). This amounts to 90 different configurations in total. Each configuration was tested with all 81 debate profiles focusing on target sets (see Section 5.4), and for every debate profile, the debate was repeated 10 times.¹⁰ So in total we have $9 \times 10 \times 81 \times 10 = 72900$ debates focusing on target sets.

The next histogram summarizes the percentage of agreement between the debates' results and the majority results for each combination of *PRO/CON* agents.

As it can be seen, when a group contains the vast majority of the agents (8 or 9 out of 10), the debate's result almost always agrees with the majority result. This is the reason why we filtered out cases of near-unanimity, and we kept only the configurations where the combination of *PRO/CON* agents was 3/7, 4/6, 5/5, 6/4 or 7/3. As a result, we focused on 50 configurations. As previously stated, all 81 debate profiles were tested for each configuration, and for every debate profile, the debate was repeated 10 times.

Another interesting element of the histogram is that the column of 7/3 (resp. 6/4) is bigger than the column of 3/7 (resp. 4/6). Furthermore, the column of 5/5 is also relatively big (in most debates, PRO wins¹¹). Apparently, the random configurations for which balanced teams of agents were (randomly) generated, slightly favor the PRO group, as far as winning the debate is concerned. This is verified in what follows, and it merits a deeper study in the future.

6.3 Analysis of the Results

As said above, we shall first examine the behaviour of the random Strategy, and then we shall focus on the remaining 81 debate profiles which focus on target sets.

⁹ We chose a small number of debated attacks, as this element causes an overhead in the computations of target sets.

¹⁰ As the agents randomly choose their moves among a set of possible moves, the results of these 10 debates may still differ.

¹¹ We remind that in case of 5/5, PRO is by default considered to be the majority.




Fig. 1. Histogram showing the coincidence of the debates' results with the majority results, for configurations having different compositions of PRO/CON agents.

For the analysis of the results and the evaluation of the strategies and debate profiles, we considered three criteria:

- **Debate length:** the average number of rounds in the debate.
- **Happiness:** the percentage of coincidence between the debate's result and the majority result. Its interest is better understood from the perspective of the debate's central authority. For example, if the central authority chooses a strategy profile for both *PRO* and *CON* (eg. the same one), then it may wish to know which one would help the majority, and which one would offer more chances to the minority.
- **Rationality:** the percentage of coincidence between the result of the debate and the merged result.

We also want to find the “best” strategies, meaning the strategies which maximize a group's chances to win the debate.

The random strategy profile We begin our analysis with the random strategy profile. As far as the maximization of a group's winning chances are concerned, the random strategy profile did not fare worse than the quite simple strategy profiles 1 and 2X. The reason is that its drawback (the fact that agents playing random attacks could harm their own group), was balanced by the drawback of profiles 1 and 2X, which can “block” a group normally able to change an attack by casting two or more votes (because in 1 and 2X, the first voter will be prohibited from casting his vote).

On the other hand, we expected that the winning percentage of a group would increase, if instead of the random profile, he used the elaborated profiles 3X and 4X. This was indeed verified, as the winning percentage always increased, up to 25% in some cases (although less in others). We also conjecture that the more attacks the GB has,

the worse the results will be for the random profile, compared to 3X and 4X. It seems logical to assume that the more attacks there are on the GB, the more harmful it is for a group to randomly play attacks, some of which may backfire.

A key disadvantage of the random profile is that, if a group uses it, then the number of rounds of the debate explodes. In most cases, when one group adopted the random profile, the number of rounds increased by a factor of 10 (eg. from 25 rounds, into 250 rounds). Remember that in the profiles focusing on target sets, if an agent has no move in a target set, he plays pass. This is not the case in the random profile, where a group can play a lot of “dummy” moves before achieving its goal.

On a positive side, if a group uses the random strategy, then the percentage of agreement with the merged outcome is quite high (in almost all cases we tested that percentage was bigger than 90%). Naturally, the reason behind this, is that that group using the random profile will cast a lot more votes during the debate, and as a result, the GB will resemble more to the merged system. This was even clearer when both groups used the random profile, when that percentage went up to 97.6%.

Concluding, the fact that the number of rounds increases dramatically when a group uses the random strategy, as well as the fact that it fares worse (as far as winning the debate is concerned) than strategies 3X and 4X, lead us to not include the random profile in the following tests, where we just compare the 9 strategy profiles focusing on target sets.

Strategies based on target sets We now turn our attention to the 9 strategy profiles focusing on target sets and their corresponding 81 debate profiles.

Each of the four graphics contains information on *all* debate profiles focusing on target sets. The top left shows the percentage of PRO wins (for every profile), the top right shows the average number of rounds of the debates, the bottom left shows the percentage of agreement between the results of the debates and the merged results, and the bottom right shows the percentage of agreement between the results of the debates and the majority results.

Let us first consider the criterion of *debate length* (top-right). The lowest number of rounds is found when both agents use strategy 4. A small number of rounds is also obtained in profiles (1,4X), (4X,1) (where $X \in \{A, B, C\}$) and (1,1). For the latter, the reason is that there are cases where a group cannot vote on an attack because no single agent can change it (and thus the debate stops). For profiles (4X,4Y) the reason debates are short is that agents are not forced to play (useless) pass-moves, as they can reinforce attacks on the GB while they are winning. This is not possible with profiles (3X,3Y) which give the longest debates. Note that agents using the strategy profiles 4X have incentive to give more information than with the other strategy profiles. That can be seen as a disadvantage for agents who wish to hide information. A last remark on debate length: If we concentrate only on rounds which do not contain pass moves (let us call them *no-pass rounds*), then the results of strategy profiles 3X and 4X are inverted. Strategy profiles 4X lead to more no-pass rounds, than profiles 3X (eg. (4C,4C) leads in average to 11.97 no-pass rounds, while (3C,3C) leads in average to 10.36 no-pass rounds). We clearly see that when profiles 3X are used, many rounds involve pass moves, and this is the reason why profiles 3X have the biggest total number of rounds.




Fig. 2. Top-Left: Percentage of wins by PRO. Top-Right: Number of rounds of the debates. Bottom-Left: Percentage of agreement with the merged result. Bottom-Right: Percentage of agreement with the majority result. PRO strategies are shown on the left side, and CON strategies on the right side of every graphic.

Let us now focus on *rationality* (bottom-left). The most “rational” outcomes (closer to the results of the merged system) are obtained when both groups use one of the strategies: 3A, 3B, 3C, 4A, 4B, 4C (the percentage of agreement being 0.88). The only cases where the results of the debates are farther from the merged results are when a group uses strategy profile 4X and the other group uses strategy profile 1 or 2X. So, we pull away from the merged result when a group uses the most advanced strategy (4X), while the other a simple one (1 or 2X). The smallest agreement is 0.66, at profile (1,4X).

Similar results are obtained when we focus on *happiness* (bottom-right). Almost all profiles give a similar value of agreement with the majority (about 0.85). However, when PRO uses strategies 1, or 2X, and CON uses 4X, the debate’s result starts to move away from the majority’s opinion (its minimum value is 0.7).

Regarding the strategy which is most likely to win a debate, the most elaborated strategies 3 and 4 provide a clear advantage. PRO’s best chance to win is when the profile (4X,1) is used (0.75 percentage of PRO winning). Similarly, CON’s best chance to win is in profile (1,4X) (0.38 of PRO winning). In general, no matter what strategy a group is using, the other group increases its winning percentage if it uses strategy 3 or 4, instead of the simpler 2 and 1 (1 being the worst choice). It is also quite clear, as mentioned before, that PRO win more debates than CON, something apparently related to the nature of the randomly generated master systems from which balanced PRO/CON groups are generated.

Finally, some remarks on the heuristics. Heuristics C which focuses on the smallest target sets, and prefers moves able to add/remove an attack, was expected to lead to the quickest debates. This was verified, although its results were not significantly better than the results of the simpler heuristics B and A. For example, the debate profile (4C,4C) lead to 23.88 rounds in average, while the profile (4A,4A) lead to 24.81. Also, the debate profile (3C,3C) lead to 35.29 rounds in average, while the profile (3A,3A) lead to 36.29. We conjecture that, when heuristics C is used instead of B or A, the decrease in the number of rounds is small, due to the fact that the randomly generated systems do not contain many target sets, and these target sets do not have great differences in size. We expect that in the case of master systems with target sets of considerably different sizes, heuristics C will lead to a more significant decrease in the number of rounds, compared to heuristics B and A.

To conclude, a general observation is that the more sophisticated strategy profiles (3X and 4X) are the best choices for the agents who want to win the debate. Their main difference lies on the average number of rounds, and on the amount of information disclosed during the debate. Surprisingly, the simpler strategy profiles (1 and 2X) offer an interesting alternative, provided that the debate's central authority can ensure that both groups will use a simple strategy profile, and that no group will switch into using a sophisticated one. It is worth noting that, in the above experiments, the probability that the winner is the same, when either profile (1,1) or profile (3C,3C) is used, was almost 95%. Finally, the use of heuristics C shortens the length of the debates, though more tests are needed in order to evaluate its impact.

7 Conclusion

We have presented a framework, where debating agents vote on attacks, focusing on a single argument. The agents' relevant expertise plays an important role on the aggregation of the votes. Some interesting properties of target sets, presented in [12], motivated us to define debate strategies focusing on them. A number of strategies and heuristics (of varying complexity) were proposed. We performed a number of experiments and drew conclusions on the strategies, using as criteria the probability of winning, the debate's length, its rationality and the agents' happiness. We also verified our intuition on the best strategies and studied the heuristics' contribution. There are many interesting directions for future research: the relation between the master system (number of arguments, debated attacks, cycles) and the debate's results, as well as the relation between different agent group compositions (eg. with low or high intra-group similarity) and the debate's results. Also, studying debates where the agents' systems may change during the debate looks promising, but challenging.

References

1. Scheuer, O., Loll, F., Pinkwart, N., McLaren, B.: Computer-supported argumentation: A review of the state of the art. *International Journal of Computer-Supported Collaborative Learning* **5** (2010) 43–102
2. Toni, F., Torroni, P.: Bottom-up argumentation. In: *Proc. of TAFA'11*. (2011) 249–262

3. Dunne, P.E., Hunter, A., McBurney, P., Parsons, S., Wooldridge, M.: Weighted argument systems: Basic definitions, algorithms, and complexity results. *Artif. Intell.* **175**(2) (2011) 457–486
4. Kontarinis, D., Bonzon, E., Maudet, N., Moraitis, P.: Picking the right expert to make a debate uncontroversial. In: *Proc. of COMMA'12*. (2012) 486–497
5. Dung, P.M.: On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* **77**(2) (1995) 321–358
6. Coste-Marquis, S., Konieczny, S., Marquis, P., Ouali, M.: Weighted attacks in argumentation frameworks. In: *Proc. of KR'12*. (2012)
7. Leite, J., Martins, J.: Social abstract argumentation. In: *Proc. of IJCAI'11*. (2011) 2287–2292
8. Egilmez, S., Martins, J., Leite, J.: Extending social abstract argumentation with votes on attacks. In: *proc. of TFAFA'13*. (2013) 16–31
9. Bonzon, E., Maudet, N.: On the outcomes of multiparty persuasion. In: *Proc. of AAMAS'11*. (May 2011) 47–54
10. Prakken, H.: Coherence and flexibility in dialogue games for argumentation. *Journal of Logic and Computation* **15** (2005) 347–376
11. Boella, G., Gabbay, D., Perotti, A., van der Torre, L., Villata, S.: Conditional labelling for abstract argumentation. In: *Proc. of TFAFA'11*. (2011)
12. Kontarinis, D., Bonzon, E., Maudet, N., Moraitis, P.: On the use of target sets for move selection in multi-agent debates. In: *Proc. of ECAI'14*. (2014) To appear.
13. Rahwan, I., Larson, K., Tohmé, F.: A characterisation of strategy-proofness for grounded argumentation semantics. In: *Proc. of IJCAI'09*. (2009) 251–256
14. Caminada, M.: Truth, lies and bullshit; distinguishing classes of dishonesty. In: *Proc. of SS@IJCAI'09*. (2009) 39–50
15. Sakama, C.: Dishonest arguments in debate games. In: *Proc. of COMMA'12*. (2012) 177–184
16. Coste-Marquis, S., Devred, C., Konieczny, S., Lagasquie-Schiex, M.C., Marquis, P.: On the Merging of Dung's Argumentation Systems. *Artificial Intelligence* **171** (2007) 740–753