King's Research Portal

[Link to publication record in King's Research Portal](#)

A heuristic strategy for persuasion dialogues

Josh MURPHY, Elizabeth BLACK and Michael LUCK

Department of Informatics, King's College London, UK

Abstract Argument-based persuasion dialogues provide an effective mechanism for agents to communicate their beliefs, and their reasons for those beliefs, in order to convince another agent of some topic argument. In such dialogues, the persuader has strategic considerations, and must decide which of its known arguments should be asserted, and the order in which they should be asserted. Recent works consider mechanisms for determining an optimal strategy for persuading the responder. However, computing such strategies is expensive, swiftly becoming impractical as the number of arguments increases. In response, we present a strategy that uses heuristic information of the domain arguments and can be computed with high numbers of arguments. Our results show that not only is the heuristic strategy fast to compute, it also performs significantly better than a random strategy.

Keywords. Argument-based dialogues, dialogue strategies, persuasion dialogues

1. Introduction

Argument-based dialogues are a useful mechanism for agent co-ordination, particularly in the domains of human-machine interaction and agreement technologies [6]. In this paper, we focus on a simple type of persuasion dialogue (where one agent presents arguments to another with the aim of convincing it to accept some argument that is the topic of the dialogue) and consider the problem of how the persuader can determine which arguments to present during the dialogue, *i.e.,* what dialogue *strategy* it should employ.

The development of methods for generating agent dialogue strategies is an active area of research [9]. So far, work on this problem has shown that computing an optimal strategy for two-party dialogues is computationally expensive, and becomes intractable as the number of arguments in the dialogue domain increases. Black *et al.* [1] consider the same simple persuasion dialogue setting that we focus on here, modelling it as a planning problem so that a planner can be used to generate an optimal strategy for the persuader, while Hadoux *et al.* [4] and Rienstra *et al.* [8] each support richer models of argument dialogue, generating optimal strategies using Mixed Observability Markov Decision Problems (MOMDPs) and a variant of the minimax algorithm respectively. While each of these approaches [1,4,8] determines an optimal strategy for the persuader, none have been shown to scale to domains with more than 10 arguments.

The key contribution of this paper is a heuristic strategy for persuasion that can easily scale to domains with 50 arguments (with computation time of less than 1 second). Although this heuristic strategy is not optimal, it gives a reasonable chance of successful persuasion and significantly outperforms a strategy that randomly selects arguments. Our heuristic strategy does not require the persuading agent to have any knowledge of the persuadee, relying only on arguments the persuader knows may exist in the domain, and uses a measure of distance from the topic argument to estimate the likelihood that any argument would (if asserted) affect the persuadee's perception of the topic acceptability.

The remainder of the paper is structured as follows. Section 2 provides the preliminary background on argumentation and argument dialogues, in particular the two-player simple persuasion dialogue we use as a testbed for our strategy. Section 3 introduces the heuristic used to estimate the likelihood of each argument to persuade the responder, and in Section 4 the strategy is formally defined. Section 5 details the experimental set-up, and Section 6 presents the results. Section 7 concludes with a discussion.

2. Argumentation and simple persuasion dialogues

Dung-style *argumentation frameworks* [3] are comprised of two key elements: arguments and attacks (the directed relationship between the arguments representing conflict).

Definition 1. *An **argument framework** is a tuple $AF = \langle A, R \rangle$, s.t. A is a set of arguments, and $R \subseteq A \times A$, is a set of attacks where $\langle x, y \rangle \in R$ is an attack, x to y.*

Given an argument framework, we can determine which *extensions* (sets of arguments) are rational for an agent to consider acceptable. While different extensions are based on different intuitions, a desirable property for a set of acceptable arguments is often that of *admissibility*. An argument is admissible with respect to a set of arguments S if all of its attackers are attacked by some argument in S, and no argument in S attacks an argument in S. For the rest of this paper, we consider an argument to be *acceptable* to an agent (w.r.t. an argumentation framework) if it is part of all maximal admissible sets. These criteria for acceptability are known as the *preferred sceptical semantics* (as in [3]).

Definition 2. *We define a function, $\mathsf{Acc}(\mathsf{AF})$, to return the set of acceptable arguments under the preferred sceptical semantics of the given argumentation framework AF.*

To investigate the effectiveness of the heuristic strategy we apply it to a persuasion dialogue (adapted from [1]) that has two participating agents: a *persuader* and a *responder*. The persuader's goal is to convince the responder of the dialogue topic (an argument). The responder replies truthfully as to whether it finds the topic acceptable given its (private) beliefs and the arguments asserted by the persuader. Agents engage in a dialogue under an argument framework — the *global knowledge* (all possible arguments in the domain, and the attacks between them) — from which their own personal knowledge is a subset.

Definition 3. *A **simple persuasion dialogue scenario**, under global knowledge $AF_G = \langle A_G, R_G \rangle$, is a tuple $\langle AF_P, AF_R, t \rangle$, such that:*
- *$AF_P = \langle A_P, R_P \rangle$, where $A_P \subseteq A_G$ and $R_P = R_G \cap (A_P \times A_P)$, is the persuader's initial knowledge base,*

- $AF_R = \langle A_R, R_R \rangle$, where $A_R \subseteq A_G$ and $R_R = R_G \cap (A_R \times A_R)$, is the responder's initial knowledge base, and
- $t \in A_P$, is the dialogue topic.

During the dialogue, the persuader and responder take turns to make utterances to one another; the persuader may assert arguments or choose to terminate the dialogue, while the responder makes a *yes* or *no* move, indicating whether it finds the topic acceptable. A *well-formed simple persuasion dialogue* is one in which the persuader only asserts arguments from its knowledge base and the responder replies truthfully, and that terminates once either the responder is convinced or the persuader chooses to give up.

Definition 4. *A **well-formed simple persuasion dialogue** of a simple persuasion dialogue scenario $\langle AF_P, AF_R, t \rangle$ under global knowledge $\langle A_G, R_G \rangle$, is a sequence of moves $[M_0^P, M_0^R, ..., M_n^P, M_n^R]$, such that:*
- $\forall i$ such that $0 < i < n$, $M_i^P \in A_P$,
- $M_n^P \in A_P \cup \{\texttt{terminate}\}$,
- $\forall i$ such that $0 < i < n$, $M_i^R = \texttt{no}$ and $t \notin \mathsf{Acc}(\langle A_R \cup \{M_0^P, ..., M_i^P\}, R_G \rangle)$,
- $M_n^R \in \{\texttt{yes}, \texttt{no}\}$, and
- $M_n^R = \texttt{yes}$ iff $t \in \mathsf{Acc}(\langle A_H \cup \{M_0^P, ..., M_n^P\}, R_G \rangle)$.

*A dialogue is **terminated** iff either $M_n^P = \texttt{terminate}$ or $M_n^R = \texttt{yes}$. A terminated dialogue is said to be **successful** iff $M_n^R = \texttt{yes}$, and **unsuccessful** otherwise.*

Over the course of a well-formed simple persuasion dialogue, the responder has no strategic concerns, as it must reply honestly if it finds the topic acceptable. However, each turn of the persuader requires a decision as to whether an argument should be asserted, and if so, which arguments in its knowledge base should be asserted. Previous work [1] has applied automated planning techniques to find an optimal strategy for the persuader to apply, but does not scale well beyond 8 domain arguments. In Section 4 we present a heuristic strategy, and show that this can easily scale to domains with up to 50 arguments. First, however, we give the intuition on which this heuristic strategy relies.

3. Evaluating the influence of arguments

We consider the *local* topological properties of argument graphs to estimate how beneficial an argument would be if asserted. The estimate is based on the intuition that arguments topologically closer to the topic are more likely to affect its acceptability. We estimate the likelihood that an argument affects the acceptability of the topic and whether the argument defends or attacks (perhaps indirectly) the topic. Note that argument acceptability not only depends on the attackers of the argument, but on the acceptability of the attackers. Thus, we are interested in *argument paths* terminating in the topic argument.

Definition 5. *An **argument path**, in an argument graph $AF = \langle A, R \rangle$ with topic t, is a list of arguments $p = [a_0, a_1, ..., a_k]$, such that:*
- $a_0 = t$,
- $\forall i$ such that $1 \leq i < k$, $\langle a_{i+1}, a_i \rangle \in R$,
- $\forall i, j$ such that $0 \leq i, j \leq k$, $a_i = a_j$ iff $i = j$ (arguments are distinct).

*The **depth** of an argument a in an argument path $p = [a_0, a_1, ..., a_i]$ is given by the function: $\mathsf{depth}(a, p) = x$ where $a = a_x$.*

Figure 1. An example argument graph with 8 arguments.

Example 1. *Consider the example argumentation framework in Figure 1 and the topic being T. Valid argument paths include $[T, F, G]$, $[T, A, B]$, and $[T, A, B, C]$; invalid argument paths include $[A, B, C]$ (the first argument is not the topic), and $[T, A, F]$ (there is no such path is in the argumentation framework).*

The distance of an argument from the topic argument provides an estimate of how likely it is that asserting the argument will affect the acceptability of the topic. The intuition behind this is as follows: for an argument to affect the topic through a particular argument path, all preceding arguments on that path must be present; furthermore, any arguments that precede the argument in question and support the topic cannot be defeated by an acceptable argument from another path. The more arguments that proceed the argument on a particular path, the more chance that one of these conditions may not hold, thus the more likely it is that the argument will not affect the topic through that path.

Example 2. *Consider the example argumentation framework in Figure 1. The persuader wishes to convince the responder (whose arguments are unknown) that the topic T is acceptable. Consider that the persuader chooses to assert the argument G; in order for this to have a chance of changing the responder's perception of the acceptability of the topic, the responder must know F. Consider instead that the persuader chooses to assert the argument D (which is twice as far away from the topic as G); for this to have a chance of changing the responder's perception of the acceptability of T, not only must the responder know A, B and C, but it must also be that the responder cannot know E.*

To obtain an estimate of how likely each argument is to affect the acceptability of the topic, we must consider all argument paths in the argument graph that start with the topic. The number of possible argument paths grows exponentially as the size of argumentation framework increases, so we consider only argument paths up to a specified depth.

Definition 6. *The **complete set of argument paths** with depth d of an argumentation framework AF and topic argument t, is a set of argument paths $C^d_{AF,t}$ where*
$$C^d_{AF,t} = \{[t, a_1, ...a_x] \mid [t, a_1, ...a_x] \text{ is an argument path in } AF, \ x \leq d, \text{ and }$$
$$\nexists [t, a_1, ...a_x, ..., a_y] \text{ such that } [t, a_1, ...a_x, ..., a_y] \text{ is an}$$
$$\text{argument path in } AF \text{ and } x < y \leq d\}$$

An argument at an even depth in a path will be a *supporting argument* of the topic, and its presence in an agent's knowledge *increases* the likelihood that it finds the topic acceptable (the argument is either the topic argument itself, or an argument that attacks an argument that attacks an opposing argument). Similarly, an argument at an odd depth will be an *opposing argument*, and its presence *decreases* the likelihood that it finds the topic to be acceptable (the argument is an attacker of a supporting argument).

With respect to a particular argument path, the magnitude of an argument's *value* is an estimation of the likelihood that the argument will affect the acceptability of the topic, and the sign indicates whether it is likely to make the topic acceptable or unacceptable.

Definition 7. *The **value** of an argument a with depth $d = \mathsf{depth}(a, p)$ w.r.t. an argument path $p = [a_0, a_1, ..., a_i]$ is given by the function:*

$$\mathsf{value}(a, p) = \begin{cases} 0 & \text{if } a \notin \{a_0, ..., a_i\} \\ 1/2^d & \text{if } a \in \{a_0, ..., a_i\} \text{ and } d \mod 2 = 0 \\ -1/2^d & \text{if } a \in \{a_0, ..., a_i\} \text{ and } d \mod 2 = 1 \end{cases}$$

To determine the *estimated utility* of an argument, we sum the values of that argument with respect to each argument path to the topic.

Definition 8. *The **estimated utility** of an argument A in an argumentation framework AF with topic t to a depth d, is a real number given by the function eu such that:*

$$\mathsf{eu}(A, C_{AF,t}^d) = \sum_{p \in C_{AF,t}^d} \mathsf{value}(a, p).$$

4. Heuristic strategy

A persuader using the heuristic strategy will not give up trying to convince the responder until it has run out of arguments to assert (known as an *exhaustive persuader* [2]). It uses estimated utility to determine which argument to assert, choosing one not yet asserted.

Definition 9. *Consider a persuader with a knowledge base $AF_P = \langle A_P, R_P \rangle$ participating in a dialogue $D = [M_0^P, M_0^R, ..., M_n^P, M_n^R]$, under a global knowledge $AF_G = \langle A_G, R_G \rangle$. The **heuristic strategy** for a depth d is given by the function $\mathsf{hStrategy}_d$ such that:*

- *if $A_P - \{M_0^P, ..., M_n^P\} = \emptyset$ then $\mathsf{hStrategy}_d(D) = \mathtt{terminate}$, otherwise*
- *$\mathsf{hStrategy}_d(D) = M$ where $M \in \{A \in A_P - \{M_0^P, ..., M_n^P\}\ |$*
  $\forall B \in A_P - \{M_0^P, ..., M_n^P\}, \mathsf{eu}(A, C_{AF_G,t}^d) \geq \mathsf{eu}(B, C_{AF_G,t}^d)\}$

Note that a persuader using the heuristic strategy can only assert arguments from its knowledge base, but uses global knowledge to determine which argument to assert. Similar to the virtual argument approach taken by Rienstra *et al.* [8], we assume that the persuader can only assert arguments it is aware of, but is aware of the potential existence of all arguments in the domain, even those that it cannot itself assert. Other works that determine strategies for argument dialogues make similar assumptions and further assume that the persuader has a model of its opponent's knowledge [1] or behaviour [4].

5. Implementation

To evaluate our heuristic strategy we generate random simple persuasion dialogue scenarios, in which the persuader selects which arguments to assert. As a benchmark for evaluation, we use a random strategy, by which a persuader will assert its unasserted arguments at random until the responder is persuaded or there are no unasserted arguments.

To generate a random simple persuasion dialogue scenario, an argument graph representing the global knowledge must be selected. In our experiments, we randomly generate two types of argument graph: tree-like and grid-like (full details of their generation are available at `github.com/joshlmurphy`). This allows us to generate a large number of dialogue scenarios on which to run experiments. Except where noted, we use

Figure 2. Percentage success rate of strategies. Error bars indicate standard error.

sparse, fully-connected, tree-like graphs which rarely contain cycles. These properties are based loosely on argument frameworks transcribed from BBC Radio 4's Moral Maze program, in which experts aim to persuade a panel of an opinion [7].

Once the global knowledge has been generated, arguments are evenly distributed into the persuader's responder's knowledge bases at random. The topic argument of the dialogue is then selected randomly from the persuader's knowledge base so that the topic is initially known by the persuader, but not by the responder. For our experiments the heuristic strategy considers argument paths up to depth 5; initial testing showed this allowed for a strong success rate while remaining fast to compute. We leave an analysis of how depth affects success strategy and computation time for future work.

The implementation for the generation and testing of simple persuasion dialogues was done in Java, and run on a standard PC (1.86 GHz dual-core processor, 2GB RAM). We used libraries from Tweety [10] to determine whether the argument topic was acceptable under the preferred sceptical semantics for a given argument graph.

6. Results

The heuristic strategy has a high success rate It is desirable for a dialogue strategy to have a high success rate in achieving an agent's dialogue goals no matter what the agents know. For simple persuasion dialogues, this means that the persuader's strategy should have a high probability of persuading the responder of the topic argument. Both the heuristic and random strategies were run on domains with 8 arguments, with different rates of argument subsets making the topic acceptable. The probability of persuader success for the strategies was determined by running many simulations of dialogues, each with a different randomly generated argumentation framework, and recording the percentage of argument subsets that make the topic acceptable in the argumentation framework, as well as whether the persuader is successful when using the heuristic or random strategy. The results are shown in Figure 2. We observe a similar trend for both strategies: as the proportion of argument subsets of the global knowledge that make the topic acceptable increases, so does the likelihood that the strategy is successful. The results show that the heuristic strategy is more likely to be successful than the random strategy.

Table 1. Time to compute heuristic strategy (seconds). *Args* is the number of arguments in the domain.

| Args	10	20	30	40	50
Time	<0.1	0.21	0.37	0.56	0.77

Figure 3. The heuristic strategy remains successful with increasing numbers of arguments.

The heuristic strategy is fast to compute   To determine the computational cost of generating the heuristic dialogue strategy, we measure the time taken to compute the estimated utility of each argument that is assigned to the persuader in a randomly generated dialogue scenario. The results are shown in Table 1, giving the average time for 1,000 random dialogue scenarios. For domains with fewer than 10 arguments the generation of the strategy took less than 0.1 seconds. At 11 arguments, the increase in time is noticeable, and appears to be somewhat linear, allowing computation of the heuristic strategy in less than a second for as many as 50 arguments in the domain. The results show that the heuristic strategy is efficiently scalable for domains with large numbers of arguments.

The heuristic strategy succeeds with many arguments   As can be seen from the results in Figure 2, the chance of successfully convincing the responder depends heavily on the particular argument graph that determines the global knowledge. The more subsets of arguments from the global knowledge that determine the topic to be acceptable, the more chance of reaching a point in the dialogue where such a set of arguments is available to the responder, causing it to terminate the dialogue successfully. To investigate how the performance of the heuristic strategy scales with the number of arguments we needed to generate global knowledge argument graphs in such a way that the proportion of argument subsets that determine the topic to be acceptable remains near constant as the size of the graphs increases. Thus, here we used partial grids, which allowed us to keep the average percentage of subsets of the global knowledge that make the topic acceptable within the range 28%–33% for all argument graphs we experimented with. We observe in Figure 3 that there is a slight decrease in the success rate of the heuristic strategy as the number of arguments increases because, as the argument graph grows, so does its complexity, and these complexities are ignored by the heuristic strategy. The decrease in success can be considered a necessary sacrifice for a computationally tractable strategy.

7. Discussion

In this paper we have presented and evaluated a heuristic strategy that can be used in persuasion dialogues. Our results show that this heuristic strategy is fast to compute,

even for domains with a large number of arguments, which is not the case for existing approaches that generate optimal strategies [1,4,8].

In future work, we intend to investigate the performance of the heuristic strategy in more complex scenarios, specifically persuasion dialogues involving more than two participants, each of which may assert arguments with the aim of convincing the others. We expect that existing approaches for determining optimal strategies [1,4,8] would be intractable here, since the probabilistic information about the opponent used determines the state space that must be searched to find an optimal solution and so as the number of opponents increases, the number of possible states to consider increases exponentially.

Argument strategies that use heuristic information have also been investigated in different types of dialogue. Kontranis *et al.* evaluate a set of heuristic-style strategies that agents use in a dialogue-type scenario, in which participants vote on the attacks between globally known arguments, with the goal to reach a consensus [5]. In comparison, the heuristic strategy we present is based on a typical dialogue game in which agents assert arguments, rather than the focus of communication being on attack relations. Wardeh *et al.* investigate PADUA, a dialogue protocol allowing agents to classify objects based on evidence from previous examples of object classification [11]. Depending on whether the opponent is agreeable or not, the persuader can select the appropriate heuristic strategy in order to increase their success rate in deciding upon their desired classification. However, Wardeh *et al.* do not investigate the scalability of their proposed strategies.

References

[1] E. Black, A. Coles, and S. Bernardini. Automated planning of simple persuasion dialogues. In N. Bulling, L. van der Torre, S. Villata, W. Jamroga, and W. Vasconcelos, editors, *Computational Logic in Multi-Agent Systems*, volume 8624 of *LNCS*, pages 87–104. Springer, 2014.

[2] E. Black and A. Hunter. Reasons and options for updating an opponent model in persuasion dialogues. In E. Black, S. Modgil, and N. Oren, editors, *Theory and Applications of Formal Argumentation*, volume 9524 of *LNCS*, pages 21–39. Springer, 2015.

[3] P. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. In *Artificial Intelligence*, volume 77, pages 321–357. Elsevier, 1995.

[4] E. Hadoux, A. Beynier, N. Maudet, P. Weng, and A. Hunter. Optimization of probabilistic argumentation with Markov decision models. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence*, pages 2004–2010. AAAI Press, 2015.

[5] D. Kontarinis, E. Bonzon, N. Maudet, and P. Moraitis. Empirical evaluation of strategies for multiparty argumentative debates. In N. Bulling, L. van der Torre, S. Villata, W. Jamroga, and W. Vasconcelos, editors, *Computational Logic in Multi-Agent Systems*, volume 8624 of *LNCS*, pages 105–122. Springer, 2014.

[6] S. Modgil, F. Toni, F. Bex, I. Bratko, C. Chesñevar, W. Dvořák, M. Falappa, et al. The added value of argumentation. In S. Ossowski, editor, *Agreement Technologies*, pages 357–403. Springer, 2013.

[7] C. Reed. Argument corpora. Technical report, University of Dundee Technical Report, Available online at `www.arg.dundee.ac.uk/corpora`, 2013.

[8] T. Rienstra, M. Thimm, and N. Oren. Opponent models with uncertainty for strategic argumentation. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence*, pages 332–338, 2013.

[9] M. Thimm. Strategic argumentation in multi-agent systems. *Künstliche Intelligenz*, 28:159–168, 2014.

[10] M. Thimm. Tweety - A comprehensive collection of Java libraries for logical aspects of artificial intelligence and knowledge representation. In *Proceedings of the Fourteenth International Conference on Principles of Knowledge Representation and Reasoning*, 2014.

[11] M. Wardeh, T. Bench-Capon, and F. Coenen. PADUA protocol: Strategies and tactics. In K. Mellouli, editor, *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*, volume 4724 of *LNCS*, pages 465–476. Springer, 2007.