A heuristic strategy for persuasion dialogues

Josh Murphy, Elizabeth Black, Michael Luck
King's College London

Introduction

Argument-based dialogues allow agents to effectively communicate both their beliefs and the reasons they have for holding those beliefs; consequently, they have become a useful mechanism for agent co-ordination, particularly in the domains of human-machine interaction and agreement technologies (Modgil et al. 2013). 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 argument to present at each step of 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 in argumentation (Thimm 2014). So far, work that has investigated 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.*'s approach (Black, Coles, and Bernardini 2014) considers 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 the approaches of Hadoux *et al.* (Hadoux et al. 2015) and of Rienstra *et al.* (Rienstra, Thimm, and Oren 2013) 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 (Black, Coles, and Bernardini 2014; Hadoux et al. 2015; Rienstra, Thimm, and Oren 2013) determines an optimal strategy for the persuader, none have been shown to scale to domains with more than 10 arguments.

We present a heuristic strategy for persuasion that we show can easily scale to domains with 50 arguments. A real-world example of such a domain is Decide Madrid[1], an online forum in which citizens can participate in debates in order to make meaningful decisions about local government policy; debates on this site commonly have in excess of 50 arguments. Although the heuristic strategy we present is not optimal, we show that it gives a reasonable chance of successful persuasion and significantly outperforms a strategy that randomly selects arguments to assert. Furthermore, our heuristic strategy does not require the persuading agent to have any knowledge of the persuadee, relying only on the arguments that the persuader knows may exist in the domain, and uses a measure of distance from the topic argument to estimate the likelihood that any particular argument would (if asserted) affect the persuadee's perception of the acceptability of the topic argument.

The heuristic used by the strategy is an estimation of how likely it is that asserting an argument will convince the persuadee that the topic argument is acceptable. The heuristic considers the *local* topological properties of the argument graph of the domain to determine some estimate of how beneficial an argument would be if asserted. The intuition of the heuristic is that arguments closer to the topic argument are more likely to affect the acceptability of the topic.

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); in order 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 the case that the responder cannot know E.

The heuristic strategy was empirically evaluated by analysing the results of simulated persuasion dialogues. Our results show that not only is the heuristic strategy fast to compute (in domains of 50 arguments the strategy takes less than 1 second to compute), it also retains a significantly better performance than a random strategy.

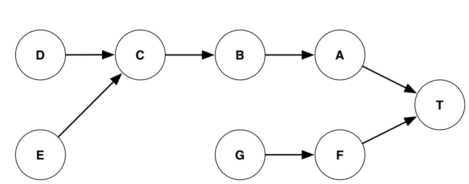

Figure 1: An example argument graph with 8 arguments.

[1] decide.madrid.es

Discussion

The heuristic strategy was evaluated by applying it to simple persuasion dialogues, in which the responder acts truthfully, and only in response to the persuader. The scenario we investigate clearly has some application to real-world scenarios: Consider the example of an agent trying to persuade an administrator to grant them privileged security permissions; the agent can assert arguments in order to convince the administrator that it should be granted, but the administrator does not have the resources to respond to all requests with any more than a notification of acceptance or rejection. The performance of our heuristic strategy depends on the arguments in the responder's knowledge base, which are unknown to the persuader. We thus expect to see similar performance of our strategy in any dialogue setting where the outcome is partly determined by arguments that are not within the persuader's control, such as more complex persuasion settings where each agent asserts arguments with the aim of persuading the other.

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 (Black, Coles, and Bernardini 2014; Hadoux et al. 2015; Rienstra, Thimm, and Oren 2013) would be intractable in such a scenario, since the probabilistic information about the opponent used by such approaches 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.

The heuristic strategy exploits knowledge that the persuader has of the dialogue domain; we make the assumption that the persuader knows about the existence of all the arguments in the global knowledge as *virtual arguments* (Hadoux et al. 2015) (the persuader is aware of the potential existence of the arguments, but may not know them enough to be able to assert them). The persuader uses this knowledge when evaluating which arguments should be asserted next. While this is a restrictive assumption, for a persuader to be effective it must have at least some knowledge either of the domain arguments, or of the persuadee's arguments. In comparison to the approach presented in this paper, other mechanisms for generating dialogue strategies are similarly restrictive in that they assume the persuader has a model of the persuadee's arguments (Black, Coles, and Bernardini 2014; Rienstra, Thimm, and Oren 2013) or of its expected behaviour (Hadoux et al. 2015). Such a model can then be updated in different ways as the dialogue progresses, and more knowledge of the persuadee is revealed (Black and Hunter 2015). However, in some domains it may be unrealistic to assume that the persuader has prior knowledge of the responder. We might expect agents to have no knowledge about the beliefs of the responder if this is the persuader's first dialogue with the responder, or if the persuader cannot be assumed to be similar to the responder. We predict the heuristic strategy's success rate could be improved by incorporating knowledge of the opponent's arguments into the utility calculation for arguments at a slight cost to computation time, in domains where such knowledge is available.

One of the challenges of empirical evaluation of argument-based dialogues is that there is a lack of benchmark domain scenarios for evaluating dialogues, and no standardised classes of argument frameworks. It is unclear what classes of frameworks exist in real-world domains, and therefore which frameworks should be used in empirical evaluations. The choice of framework, including the initial distribution mechanism of arguments to participants, appears to have a significant effect on the resulting dialogue (Murphy, Black, and Luck 2015). Future work will investigate how the structure of the underlying argumentation framework affects argument dialogues, as well as other argument-based systems. We will consider argumentation framework structures from argumentation domains (such as frameworks generated from different argument schemes (Walton, Reed, and Macagno 2008)), as well as topological structures of framework (such as trees and grids).

Acknowledgements Josh Murphy is a PhD student funded by EPSRC.

References

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

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

Hadoux, E.; Beynier, A.; Maudet, N.; Weng, P.; and Hunter, A. 2015. Optimization of probabilistic argumentation with Markov decision models. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2004–2010. AAAI Press.

Modgil, S.; Toni, F.; Bex, F.; Bratko, I.; Chesñevar, C. I.; Dvořák, W.; Falappa, M. A.; Fan, X.; Gaggl, S. A.; García, A. J.; et al. 2013. The added value of argumentation. In Ossowski, S., ed., *Agreement Technologies*. Springer. 357–403.

Murphy, J.; Black, E.; and Luck, M. 2015. Arguing from similar positions: an empirical analysis. In *Proceedings of the 3rd International Workshop on the Theory and Applications of Formal Argumentation*, volume 9524 of *LNCS*, 177–193. Springer.

Rienstra, T.; Thimm, M.; and Oren, N. 2013. Opponent models with uncertainty for strategic argumentation. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, 332–338. AAAI Press.

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

Walton, D.; Reed, C.; and Macagno, F. 2008. *Argumentation schemes*. Cambridge University Press.