An Automated Planning Approach for Generating Argument Dialogue Strategies

Tanja Daub, Elizabeth Black and Amanda Coles
Department of Informatics, King's College London, UK

Abstract

Argumentation dialogues are a well established method of solving conflict in multi-agent systems. In recent years, different ways of generating effective dialogue strategies, which determine the arguments an agent should assert, have been investigated. We approach persuasion dialogues as a classical planning problem. We would like to build on the ongoing EPSRC project *Planning an Argument* (King's College London) that currently is able to find simple plans, i.e. predetermined sequences of moves irrespective of the opponent's moves. We would like to find several simple plans and then merge them into a policy to generate strategies that take the opponent's moves into consideration. We hope that this approach will scale well to larger problems.

Introduction

In multi-agent settings, conflict can arise if agents have different view points. One way to resolve such conflicts is to engage in a persuasion dialogue. Typically such dialogues are based on an argumentation framework, which is defined as a set of arguments and an attack relation between those arguments (Dung 1995). Each dialogue participant has a knowledge base (a set of arguments known to the agent) and they make moves (utterances) to one another. A dialogue terminates when one participant persuades the other, or when the participants give up trying to persuade each other.

In recent years, research in this area has been concerned with finding strategies for persuasion dialogues that determine which argument the proponent should assert to its opponent, aiming to maxmise the probability of success. Identifying such an optimal strategy is difficult if the opponent's knowledge base is not known to the proponent and if there exists no knowledge about their expected behaviour. The proponent may also put themselves at a disadvantage by asserting an argument that may be useful to the opponent.

In our work, we assume that the proponent has some uncertain knowledge about the opponent's beliefs, represented as an opponent model (i.e. a probabilistic model of the possible sets of beliefs that the opponent may hold) and assume no knowledge about the opponent's strategy.

Related Work

Recent research has considered how to generate strategies for argument dialogues. Rienstra, Thimm, and Oren (2013) assume that the proponent possesses some uncertain knowledge of the opponent's beliefs and apply a variant of the minimax algorithm to generate an effective strategy. They run experiments on domains with 10 arguments to investigate the effectiveness of their approach under different levels of uncertainty around the opponent's beliefs. Black, Coles, and Bernardini (2014) apply automated planning to simple asymmetrical persuasion dialogues where the opponent does not assert any arguments and only states truthfully after each proponent move whether they now accept the topic or not. They show that it is possible for an automated planner to find optimal solutions, i.e. plans that maximise the probability of success given the proponent's uncertain knowledge of the opponent's beliefs, for domains with up to 9 arguments. Finally, Hadoux et al. (2015) apply Mixed Observable Markov Decision Processes to generate optimal policies for persuasion. They do not assume any knowledge of the opponent's initial belief state but do assume probabilistic knowledge of the opponent's expected behaviour. They show that their approach scales to domains with up to 8 arguments and acknowledge that further improvements are likely needed for it to handle examples with more than 12 arguments.

Each of these approaches assumes knowledge of the opponent's strategy. The adaption of the minimax algorithm applied by Rienstra et al. (2013) assumes that the opponent will behave optimally; the simple persuasion dialogues considered by Black et al. (2014) assume that the opponent's behaviour is deterministic; Hadoux et al. (2015) assume the proponent has a probabilistic model of the arguments the opponent will assert during the dialogue. In contrast to the research discussed above, we are interested in a scenario where no knowledge of the opponent's behaviour is available, but the proponent does have an uncertain opponent model that assigns probabilities to sets of beliefs representing the likelihood that they are known to the opponent at the start of the dialogue. By taking a planning approach, we hope to be able to harness the advances made in developing efficient planners to be able to scale to larger problems.

Planning Strategies for Argument Dialogues

The ongoing EPSRC project *Planning an Argument* (King's College London) is investigating the use of automated planning for generating strategies for persuasion dialogues. The project's approach is based on treating argumentation dialogues as classical planning problems. A classical planning task consists of a set of state variables, a set of actions that are defined by preconditions and effects, and a start and a goal state. The state space is the directed graph of all states that can be reached by the application of actions. A planner can find a sequence of actions that leads from the start state to the goal state. In classical planning, only deterministic problems without uncertainty are considered.

In a persuasion dialogue, the utterances by each agent can be represented as moves. Goal states are any states where the opponent accepts a given topic. One challenge of representing a persuasion dialogue as a planning problem is dealing with uncertainty. Uncertainty arises due to the fact that only probabilistic knowledge of the opponent's initial state is available and that the moves the opponent will make are unknown. This can be dealt with by compiling away the uncertainty by pushing it into the state description (Albore, Palacios, and Geffner 2010) so that the states contain information about each possible dialogue that could have occurred up to that state for each possible opponent model.

The project is currently concerned with finding simple strategies, i.e. strategies that follow a predetermined sequence of moves rather than responding to arguments asserted by the opponent, which guarantee a certain probability of success regardless of the opponent's behaviour. This approach is not optimal, but the strategies it produces have a reasonable probability of success and it can currently cope with examples with up to 15 arguments. We would like to extend the space of strategies by considering policies that depend on the opponent's moves.

We intend to increase the probability of success guaranteed by a simple strategy, in the first instance, with the following approach. Consider that the proponent's uncertain model of the opponent indicates that the set of beliefs available to the opponent at the start of the dialogue is one of $M = \{\Theta_1, \ldots, \Theta_n\}$. Techniques developed in the *Planning an Argument* project allow us to find simple strategies that are guaranteed to be successful for some subset of these possible opponent models, no matter which moves the opponent chooses to make. By choosing an appropriate partition $\{M_1, \ldots, M_j\}$ of M we can find a set of simple strategies $\{S_1, \ldots, S_j\}$ such that each S_i is guaranteed to be successful for all possible opponent models captured in M_i. A sensible approach to this may be to first find the simple plan that has the highest probability of success overall, and then identifying which opponent models it is not effective for. Then we can replan for these opponent models. If we can then find a way to merge these simple strategies into a policy, we can ensure success for all possible opponent models.

There are several research questions we need to explore in order to develop this approach.

- How can the set of possible opponent models be partitioned in a meaningful way into subsets that are dealt with by one simple plan?
- How can we merge several simple strategies into a policy that accounts for the opponent's moves?
- How close is the resulting policy to being optimal? Optimality in our approach is compromised because the best optimal policy may not necessarily be built on the optimal simple strategy. It may be necessary to start from simple strategies that are not optimal.
- How does this approach compare to other approaches? In particular, how well does it scale to larger problems?

Similar research has been conducted in planning by Muise, McIlraith, and Beck (2012), who propose an approach for fully observable non-deterministic planning problems in which every non-deterministic action is replaced with deterministic ones that represent all possible outcomes. Then a weak plan, i.e. a plan that is successful if all actions have the desired outcome, is generated that will lead to a successful outcome, given selected opponent moves. A policy is built incrementally by calling a planner whenever a state is encountered for which there is no successful plan. We would like to identify possible connections to their work and find a way to utilise it.

Conclusion

We intend to use planning to find policies for persuasion dialogues by finding a simple plan that yields the highest probability of success and then identifying opponent models this plan will not succeed with. Then we want to replan using those opponent models and iteratively build a policy that covers all possible opponent models.

Acknowledgements

This work was partially supported by the UK Engineering and Physical Sciences Research Council [grant ref. EP/M01892X/1]. Tanja Daub is supported by a UK Engineering and Physical Sciences Research Council Doctoral Training Grant.

References

Albore, A.; Palacios, H.; and Geffner, H. 2010. Compiling uncertainty away in non-deterministic conformant planning. In *ECAI*, volume 215, 465–470.

Black, E.; Coles, A.; and Bernardini, S. 2014. Automated planning of simple persuasion dialogues. In *Computational Logic in Multi-Agent Systems*. Springer. 87–104.

Dung, P. M. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence* 77(2):321–357.

Hadoux, E.; Beynier, A.; Maudet, N.; Weng, P.; and Hunter, A. 2015. Optimization of probabilistic argumentation with markov decision models. In *International Joint Conference on Artificial Intelligence*, 2004–2010.

Muise, C. J.; McIlraith, S. A.; and Beck, J. C. 2012. Improved non-deterministic planning by exploiting state relevance. In *ICAPS*.

Rienstra, T.; Thimm, M.; and Oren, N. 2013. Opponent models with uncertainty for strategic argumentation. In *IJCAI*, 332–338.