LEARNING THE PROBLEM REPRESENTATION FOR IMPROVING NEGOTIATION STRATEGIES

by: Eddy Fledderus
E.A.Fledderus@student.tudelft.nl
Supervisors:
Bram Renting, Pradeep Murukannaiah

BACKGROUND

A negotiation problem consists of multiple agents trying to reach an agreement. This negotiation is done in a certain domain where each domain can be different from another, for example in terms of how cooperative the domain is. The difference in these domains is something that could be used by one of these agents to improve its strategy. For this research a machine learning based agent was used:

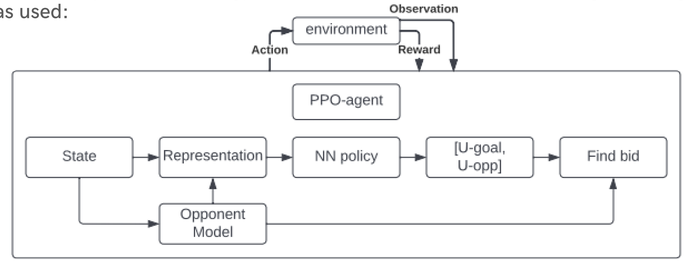

RELATED WORK

In a related research features of the domain were used to select a pre-existing algorithmic method [1]. These features were also used in this research, however, they were used in a Neural Network Policy to be directly used in the strategy. The advantage of using the features during negotiation is that you can also derive unknown features:

Features		
Common[1]	Private[1]	Derived
Amount of issues	Average utility of all proposals	Standard deviation of the opponent's bids
Amount of all values in issues	Standard deviation of all values in issues	Average utility of the opponent's bids
Amount of possible proposals	Standard deviation of the utility of all proposals	Utility of the best opponent bid
		The perceived Nash point
		The perceived improvement of opponent bids

METHODOLOGY

To get a clear view of how much of an effect these features have on the results, the features have been divided into 4 groups: common, private, exact- and inexact derived. These features have been added to a base agent using the following features: the utility of the last 3 bids and the progress of the negotiation. The utility of these agents negotiating against a group of agents have been compared to a base agent in the same situation. After this, an analysis has been done on how detrimental the features were in the negotiation.

RESULTS/FINDINGS

The results of the agents with the extra features were promising as they were all outperforming the base agent at some point. They also all had a higher maximum utility score over the training increments.

ANALYSIS

After analyzing the results, they seemed less optimistic. In the analysis the additional issues of the agents were set to 0, effectively making the issues useless throughout the negotiation. For the common agent and the exact derived agent, this made a big difference. Thus these features did help in the negotiation. For the other two agents, however, this did not make much of a difference. The private agent even improved with this change.

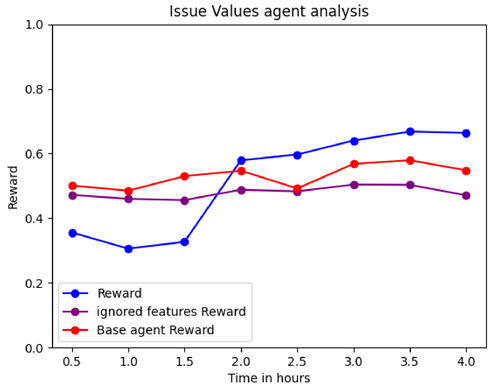

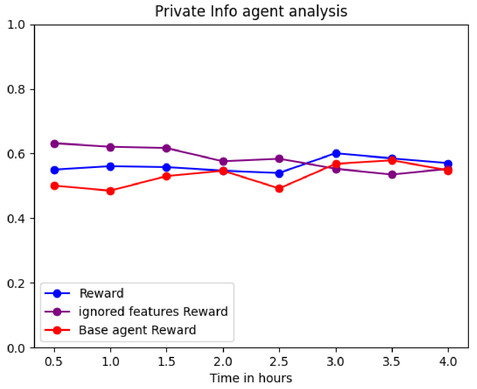

CONCLUSION

Drawing a conclusion from the analysis, it could be said that a general understanding of the size of the domain helps in obtaining a higher utility. The same can be said about having a general understanding of the opponent.

The fact that the private agent improved is something worth investigating in future research. A possible conjecture is that setting the features to 0 made the policy "think" it was in an unfavorable negotiation. This both improved personal utility and social welfare.

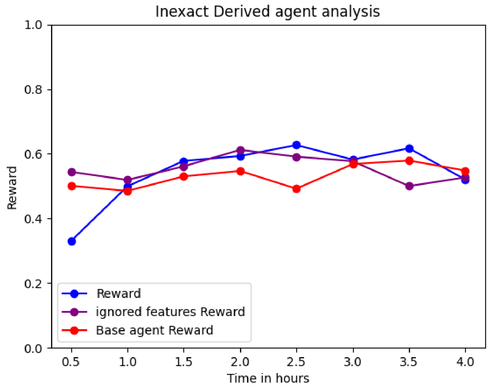

RELATED LITERATURE

[2] Ilany, L., Gal, Y. Algorithm selection in bilateral negotiation. Auton Agent Multi-Agent Syst 30, 697–723 (2016).