

AdamAgent: A Learning Moving Average Agent Submitted to The ANAC 2022 SCM League OneShot Track

Dorin Keshales Eran Hirsch

Department of Computer Science, Bar-Ilan University, Israel

June 12, 2022

1 Introduction

In 2021, a new track called SCML-OneShot was introduced. The OneShot track implements a simplified problem in which the agent can focus on the many to many concurrent negotiation problem without needing to worry about long term planning or production planning as is the case with the known standard and collusion tracks. In the SCML-OneShot an agent's utility (i.e., its profits) on a given day is completely determined by the set of contracts it secures on that day. In particular, it is independent of any past or future contracts.

As mentioned in the tutorial 'Developing an agent for SCML2021 (OneShot)'¹, simple agents do not change their behavior over simulation time, i.e. whatever they do in one day depends only on information about what is happening on that day, which may not be optimal. Therefore, the *LearningAgent* agent tries to learn about different partners over the simulation time in a simplistic manner by changing the limits of the price it is willing to accept based on accumulated statistics per agent. The main idea of the LearningAgent is to limit the minimum price (for selling) and/or maximum price (for buying) based on events not only in this day but up until now.

2 The Design of AdamAgent

2.1 Negotiation Choices

Our agent builds upon the LearningAgent, but fixes one of its downfalls. The problem with the LearningAgent is that it is hard-headed, because once it sees

¹http://www.yasserm.com/scml/scml2020docs/tutorials/02.develop_agent_scml2020.oneshot.html#oneshotagent

a good price from a negotiator, it insists on it for future negotiations. In other words, it assumes that the price will only improve over time. We conjecture that the price can also worsen, and that time is an important factor - the recent operations are more representative of the price than older actions.

Given the above, we suggest using an Exponential Moving Average (EMA, see equation 1). For this purpose, our agent tracks all the offers over the entire simulation, both per negotiator and for all negotiations. Once it needs to offer a price or respond to an offer, it will calculate the EMA for the specific negotiator and use it to decide the price range. The smoothing factor α is a hyperparameter which is fine-tuned based on past data on every negotiation.

$$EMA_i = Price(i) \times \alpha + EMA(i-1) \times (1 - \alpha) \quad (1)$$

3 Evaluation

We tested AdamAgent locally in simulations with AdaptiveAgent, LearningAgent and winners from last year. The results are not very stable ², but we saw that AdamAgent is a competitive agent.

Conclusions

Following last years' results it seems that an average of previous negotiations proves a competitive agent (2021 winner). We hope to improve on these results by adding a learning agent, taking into consideration the time in which the negotiations occurred and fine-tuning hyperparameters based on the data collected during the simulation.

²<https://github.com/yasserfarouk/scml/discussions/95>