June 10, 2022

DreamTeam109 ANL2022 Agent Strategy

Raz Vikelman, Itay Mastor, Inbal Yosef, Linoy Yosef
School of Computer Science, The College of Management Academic Studies, Israel

1 Introduction

The Automated Negotiating Agent Competition (ANAC) [1] is an international tournament that has been running since 2010 to bring together researchers from the negotiation community.
Our team took part in the Automated Negotiation League (ANL) 2022 [2], in which the task is to design a negotiation agent for bilateral negotiation that can learn from every previous encounter while the tournament progresses. The highest scoring agents based on individual utility and social welfare win.

2 Design

Basically, the purpose of our agent is to balance between two factors:

1. maximizing the chances of getting a great utility by accepting a generous offer or by our offer being accepted.
2. minimizing the chances of a breakdown in the negotiation session.

To do so, we must be familiar with a dictionary of a few concepts:

- **Progress Threshold**: a number between 0 and 1 that represents the point when we believe that if we don't accept an offer, we won't be able to receive another one before the deadline - so we will accept any offer. Calculated by the average duration of the latest turns, the session duration and a configurable (by what we learn about our opponent) multiplier.
- **Light Progress Threshold**: represents the point when we are not completely desperate yet, but we understand that the chances of getting a perfect utility aren't great so we are willing to compromise a little, offer and accept some more reasonable offers.
  Calculated with the same practice as the regular threshold, but 5 times further from the session deadline.
- **Top Bids Percentage**: represents the part of offers that we are willing to offer/accept from the very beginning of the session. Determined according to past sessions of negotiation with the opponent.

2.1 Bidding Strategy

- At the beginning, we pick one of our **Top Bids** randomly. This will go on until we get to the light progress threshold, so if we accept an offer and the progress is not close to 1, then probably our utility is.
- At the **Light Threshold,** we offer the best bid that our opponent ever offered to us, which we save through the session.
  We assume that it has a pretty good utility compared to just accepting any offer because of time pressure, and that a rational opponent is likely to accept it.

2.2 Acceptance Strategy

- We will accept an offer in one of three cases:
  - The offer is **very good** - as good as we could offer by ourselves in the beginning of the session.
  - We believe that this is our last chance to get more than nothing - if the progress is **greater than our threshold,** we will accept the offer.
  - The end is near and the offer is satisfactory - if the progress is **greater than our light threshold, AND** the offer is one of the **top 20%** bids in the pool for us, we will accept the offer.

2.3 Opponent Modeling

- As mentioned before, we keep track of our opponent's bids so when the session approaches the end, we can return to him with the offer which is **best for us** and we know that is **good for him**.

2.4 Learning method

- We save a **separate JSON file** for each one of our opponents, containing an array of objects that represent each session.
- As mentioned before, the top bids percentage and the progress thresholds are affected by configurable variables which are calculated based on **past negotiation sessions** with the same opponent:
    - To determine our thresholds, we use variables that say **how many rounds** before the end we would like to activate our "compromise mode". **Those variables start at 0**, because we wouldn't want to miss on the opportunity to get the most out of negotiation with agents that would rather accept bad offers and not take the risk of a breakdown in negotiation.
      Once we realize that we are facing a hardlining opponent which results in breakdowns, we will quickly take a more compromising approach.
    - A similar logic is applied to the **top bids percentage**, which will increase if we get low utility against an opponent, because it indicates that he couldn't offer or accept bids as good as we were looking for, so eventually we had to accept a bad offer from him at the last moment.

3 Results

- To test our agent, we arranged a tournament with 8 good agents (which beat the other agents) from the github repository, and 6 different domains.
- We made sure we include an extremely hardlining agent, so we don't forget to deal with this kind of situation.
- When playing 500ms rounds, our agent won first place with about 0.7 utility points per session.
- When increasing to 2000ms rounds, our agent improved to around 0.8 utility points per session. We believe that in a longer time we get a better chance (statistically) to reach a quality agreement, and that our progress thresholds are higher when there are more rounds in a session, so we get a chance to see our opponents compromise more before we do it ourselves.

References

[1] "ANAC2022 - 13th Automated Negotiating Agents Competition", http://web.tuat.ac.jp/~katfuji/ANAC2022/index.html
[2] "ANL2022 - Automated Negotiation League 2022", https://web.tuat.ac.jp/~katfuji/ANAC2022/genius.html