ANL 2024 Agent Submission

Ilan Brilovitch, Chen Shalev, Yehuda Daniel, Tal Teri

1 Introduction

In this report, we present the design and implementation of our negotiation agent, named The Dealmaker, for the Automated Negotiation League (ANL) 2024. The primary objective of The Dealmaker is to negotiate bilaterally with other agents efficiently, aiming to maximize its individual utility without knowledge of its opponent's reservation value. We leverage concepts from game theory and negotiation strategy to develop The Dealmaker as a competitive participant in the ANL.

2 Design Overview

The Dealmaker is designed as a bilateral negotiation agent equipped with its own utility function and access to its opponent's utility function, excluding the opponent's reservation value. The agent follows the Alternating Offers Protocol (AOP) where negotiation proceeds through offer exchanges until an agreement is reached or the negotiation fails. The Dealmaker's design is informed by game theoretic principles, focusing on strategic decision-making and utility maximization.

3 Agent Components

3.1 Bidding Strategy

Our choice of bidding strategy for The Dealmaker was the result of extensive research and experimentation aimed at identifying an approach that balances exploration and exploitation in negotiation settings. Initially, we experimented with simple bidding heuristics, but found them inadequate for achieving competitive results. Drawing inspiration from game theory, particularly the concepts of utility maximization and dynamic programming, we developed a sophisticated bidding strategy for The Dealmaker.

The strategy involves generating offers based on a rational selection process, considering both immediate utility gain and long-term negotiation dynamics. The Dealmaker evaluates potential offers using its utility function, prioritizing

offers with utility higher than its reserved value. The selection process incorporates a dynamic adjustment factor based on relative time in the negotiation session and the agent's aggressiveness parameter, allowing The Dealmaker to adapt its bidding strategy to varying opponent behaviors and negotiation contexts.

3.2 Acceptance Strategy

In designing The Dealmaker's acceptance strategy, we aimed to strike a balance between risk-taking and risk-aversion, informed by negotiation theory and empirical analysis of negotiation outcomes. We experimented with different acceptance criteria, ranging from conservative thresholds to more aggressive strategies, to identify an approach that maximizes The Dealmaker's payoff while minimizing the risk of unfavorable agreements.

The acceptance strategy evaluates received offers against an adjusted aspiration level derived from a dynamic function of relative time and the agent's reserved value. By strategically adjusting its acceptance threshold, The Dealmaker can exploit favorable offers while maintaining flexibility in negotiation outcomes. This adaptive approach enables The Dealmaker to navigate negotiation dynamics effectively, optimizing its utility gain over multiple negotiation sessions.

3.3 Reservation Value Modeling

One of the key challenges in negotiation is estimating the opponent's reservation value, as it directly influences the negotiation dynamics and outcomes. Our approach to reservation value modeling for The Dealmaker involved leveraging advanced estimation techniques grounded in game theoretic principles. We collected data on opponent offers and corresponding utilities over multiple negotiation sessions, using curve fitting algorithms to iteratively refine our estimation.

This adaptive modeling approach allows The Dealmaker to dynamically update its negotiation strategy based on evolving opponent behaviors and preferences. By accurately estimating the opponent's reservation value, The Dealmaker gains a strategic advantage in negotiation, enabling it to anticipate opponent moves and optimize its bidding and acceptance strategies accordingly.

Our journey to develop and refine The Dealmaker's negotiation strategy has been long and challenging. We started with simple heuristics and struggled to achieve competitive results, often trailing behind default bots and lower-ranking competitors in the competition. However, through perseverance and iterative improvement, we gradually enhanced The Dealmaker's performance, incorporating insights from game theory, negotiation theory, and empirical analysis to develop a robust and competitive negotiation agent.

From initially struggling to beat default bots to now consistently achieving competitive results, our journey with The Dealmaker demonstrates the transformative power of iterative improvement and adaptive learning in negotiation

settings. We believe that The Dealmaker's success in the ANL 2024 tournament is a testament to our dedication to pushing the boundaries of negotiation research and practice, and we are excited to showcase the effectiveness of our approach in addressing complex negotiation challenges.

4 Implementation Details

The Dealmaker is implemented in Python using the NegMAS framework, incorporating sophisticated algorithms and data structures to support its negotiation strategy. The agent class `TheDealMaker` inherits from the `SAONegotiator` class, providing interfaces for negotiation state handling and response generation. The Dealmaker's implementation reflects our deep understanding of game theory and negotiation dynamics, integrating theoretical insights into practical negotiation strategies.

5 Evaluation Strategy

During the ANL 2024 tournament, The Dealmaker will be evaluated based on its individual utility performance, reflecting its ability to achieve favorable outcomes in bilateral negotiations. The agent's strategic decision-making, informed by game theoretic principles, enables it to adapt dynamically to various opponent behaviors and negotiation scenarios, maximizing its average utility across multiple negotiation sessions.

6 Conclusion

The Dealmaker represents our commitment to leveraging game theory and negotiation theory to develop competitive negotiation agents capable of maximizing individual utility in bilateral negotiations. By integrating theoretical insights with practical implementation, The Dealmaker demonstrates the potential of strategic decision-making and adaptive learning in negotiation settings. We believe that The Dealmaker's success in the ANL 2024 tournament is a testament to our dedication to pushing the boundaries of negotiation research and practice, and we are excited to showcase the effectiveness of our approach in addressing complex negotiation challenges.