CARC: An agent based RVFitting

Hongji Xiong, Tianzi Ma, Xuan Wang, and Yulin Wu

Harbin Institute of Technology, Shenzhen

xionghj@stu.hit.edu.cn, mtz982437365@gmail.com,
wangxuan@cs.hitsz.edu.cn, yulinwu@cs.hitsz.edu.cn

April 15, 2024

Abstract

Our team has developed a highly efficient and effective automated negotiation agent for handling bilateral negotiations. Utilizing the alternating offer protocol for this tournament, we have devised a structured workflow divided into three key components: opponent modeling, acceptance strategy, and offer strategy. This approach enables us to estimate the opponent's undisclosed reserved value through curve fitting. Our acceptance strategy involves evaluating the current offer against our anticipated offer to determine whether to accept it. Lastly, our offer strategy aims to maximize our benefits throughout the negotiation process by generating offers in descending order relative to the negotiation's timeline.

1 Introduction

Negotiations are complex interactions that require intelligent strategies to achieve optimal outcomes. Our automated agent leverages advanced computational methods to dynamically adapt to opponent behaviors, manage offers, and decide on acceptances strategically. This report details the components of our agent's strategy, including opponent modeling, acceptance strategies, and offer strategies.

2 The Design of CARCAgent

2.1 Bidding Strategy

The bidding strategy of our CARCAgent is crafted to exploit the full potential of the negotiation space, guided by Nash equilibrium and Pareto optimality concepts. We start by identifying the Nash Point $N(u_{nash}(a), u_{nash}(b))$ using the `negmas.preferences.nash_points` function, which marks a fair balance of utility between the negotiating agents.

- **Establishing Offer Range:** The range is delineated by two boundary points on the Pareto frontier. The left boundary, p_{left}, is set at $(u_{nash}(a) - x_a, u_{nash}(b) + x_b)$, making the offer more appealing to the opponent initially. The right boundary, p_{right}, trends towards the opponent's reserved value, calculated at $(u_{nash}(a) + y_a, u_{b_{reserve}})$, strategically approaching their minimum acceptable utility.
- **Dynamic Offer Adjustment:** As the negotiation progresses, the offers are dynamically adjusted from the right towards the left boundary. This strategic shift allows for gradual concessions, maintaining a competitive edge while fostering agreement prospects.

2.2 Acceptance Strategy

Our acceptance strategy integrates time-dependent dynamics to maximize utility while accounting for the negotiation timeline and opponent behavior:

- **Early Phase Strategy:** In the early stages of negotiation, our agent employs a stringent acceptance criterion, only accepting offers that significantly exceed our next projected concession.
- **Approaching Deadline:** As the deadline looms, the strategy becomes more lenient, accepting any offers that meet or exceed our reserved value. This shift is crucial to avoid breakdowns in negotiation that could result in suboptimal or zero utility.

2.3 Opponent Modeling

Opponent modeling is pivotal in understanding and predicting the behavior of the counterparty in automated negotiation systems. Our model hypothesizes a set of possible reserved values for the opponent, denoted as $H = \{H_1, H_2, \ldots, H_n\}$, based on prior knowledge or historical data. This set spans a probable range centered around our reserved value, adjusted by a variable constant to cater to different negotiation scenarios.

2.3.1 Initialization of Hypotheses

Each hypothesis H_i within the range $[rv1 - 0.1, rv1 + 0.1]$ is assigned an initial probability. These probabilities are adjusted higher for hypotheses closer to the Pareto boundary, reflecting the likelihood of the opponent's true reserved value.

2.3.2 Evidence Collection

As the negotiation progresses, we gather evidence based on the opponent's behaviors—specifically their offers and rejections:

- **Rejections:** An opponent's refusal of a highly advantageous offer suggests a higher reserved value than currently estimated, prompting an increase in the probability of higher reserved value hypotheses.
- **Offer Patterns:** Rapid concessions followed by a steadfast offer indicate proximity to the opponent's actual reserved value.

2.3.3 Updating Hypotheses

Utilizing Bayesian updating rules, we recalibrate the probabilities of each hypothesis whenever new evidence is observed. This dynamic adjustment refines our predictions about the opponent's reserved value as more data accumulates.

2.3.4 Dynamic Reserved Value Adjustment

The model selects the hypothesis with the highest posterior probability as the current estimate for the opponent's reserved value. This estimate guides our offer strategy, particularly in the early stages of negotiation. As negotiations approach the Nash equilibrium point, the frequency and magnitude of updates decrease, stabilizing the estimate of the reserved value.

2.3.5 Application in Strategy

The opponent modeling directly informs our bidding strategy, shaping initial and subsequent offers and enabling strategic concessions that align with the evolving understanding of the opponent's preferences and thresholds.

3 Evaluation

4 Lessons and Suggestions

Conclusions