NayesianNice: An agent submitted to the ANAC 2024 ANL league

Shengbo Chang
Tokyo University of Agriculture and Technology,
changshengbo1997@gmail.com.

April 7, 2024

Abstract

This report introduces the NayesianNice agent for the ANL of ANAC 2024. The agent implements a negotiation strategy that integrates a mixture time-dependent bidding strategy, opponent modeling through Bayesian modeling, and an acceptance strategy. At its core, the agent targets the predicted Nash outcome in its bidding strategy.

1 Bidding Strategy

The bidding strategy, as defined in Eq. 1, enables the agent to dynamically adjust its negotiation strategy throughout the negotiation process. Initially, the agent adopts an aggressive stance, aiming to maximize its utility at the early stages of negotiation. As the deadline approaches, the strategy shifts towards more accommodating approaches, optimizing for agreement rather than maximal gain. The strategic incorporation of random noise allows the agent to effectively respond to and potentially counter tit-for-tat strategies employed by opponents. Importantly, U_{BestNash} represents the Nash outcome assuming the opponent's reservation value is zero, $U_{\text{PredictedBest}}$ signifies the best possible outcome given the predicted opponent's reservation value (precisely, the outcome that meets the opponent's reservation value), and $U_{\text{PredictedNash}}$ denotes the Nash outcome based on the predicted opponent's reservation value.

$$
\begin{aligned}
&U_t = t \cdot asp_1 + (1-t) \cdot \max(asp_2, asp_3) + noise,\\
&\text{where}\\
&asp_1 = (U_{\max} - U_{\text{PredictedNash}}) \cdot (1-t^5) + U_{\text{PredictedNash}},\\
&asp_2 = (U_{\max} - U_{\text{PredictedBest}}) \cdot (1-t^5) + U_{\text{PredictedBest}},\\
&asp_3 = (U_{\max} - U_{\text{BestNash}}) \cdot (1-t^5) + U_{\text{BestNash}},\\
&\text{if } t < 0.8,\ \text{then } noise = \text{Uniform}(0, 0.005),\ \text{else } noise = 0.
\end{aligned}
\tag{1}
$$

2 Acceptance Strategy

Immediate Acceptance at Deadline: If the negotiation has reached its final step, the agent accepts any offer that provides a utility greater than its reserved value.

Comparison with Proposed Utility: For offers received before the final step, the agent accepts the offer if its utility is equal to or greater than the utility of what it plans to propose next.

3 Opponent Modeling

The agent utilizes a Bayesian prediction method to estimate the opponent's reservation value. This approach hypothesises potential reservation values and negotiation behaviors to update beliefs about which of these hypotheses might be true. Especially, upon receiving a new bid, our model updates the prediction from the original priors with only the latest five bids.

3.1 Hypothesis Space and Prior Distribution

The hypothesis space is constructed for both the reservation value, r, and the bidding strategy parameter, b.

The hypotheses for the opponent's reservation value are derived from a range that spans from the highest utility bid observed from the opponent, indicating a potential lower bound of their reservation value, to the highest conceivable reservation value that allows for at least one mutually acceptable negotiation outcome. We set the number of hypothesis to 10.

The hypotheses for the opponent's bidding strategy parameter are heuristically guessed to be $B = \{1, 2, 5\}$

A uniform prior distributions for both the reservation value hypotheses and the bidding strategy parameters:

$$\text{Prior}(r) = \frac{1}{N_R}, \quad \forall r \in R \tag{2}$$

$$\text{Prior}(b) = \frac{1}{N_b}, \quad \forall b \in B \tag{3}$$

where R and B are the sets of all possible hypotheses for r and b, respectively, with N_r and N_b denoting their sizes.

3.2 Likelihood Function

The likelihood of observing a specific bidding utility of the opponent u_{opp} at time t, given hypotheses r and b, is modeled as:

$$\mathcal{L}(u|r, b, t) = \exp\left(-\frac{(u_{opp} - ((1-r)\cdot(1-t^b)+r))^2}{2\sigma^2}\right) \tag{4}$$

Here, σ represents the standard deviation of the Gaussian noise, reflecting the uncertainty in our observations.

3.3 Posterior Calculation

The posterior probability for each hypothesis a is updated by integrating over all possible values of b, effectively marginalizing over the bidding strategy hypotheses:

$$\text{Posterior}(r|u,t) = \text{Prior}(r) \cdot \sum_{b \in B} (\mathcal{L}(u|r,b,t) \cdot \text{Prior}(b)) \tag{5}$$

3.4 Prediction

Upon receiving a new bid, our model updates the prediction from the original priors with only the latest five bids.

To calculate the posterior for each hypothesis, we apply the Bayesian update rule using the likelihood of the observed utilities from these recent bids. The posteriors are then normalized to ensure they sum to 1:

$$\text{Posterior}(r|u,t) \propto \text{Prior}(r) \times \text{Likelihood}(u|r,t) \tag{6}$$

$$\text{Normalization: } \sum_{r \in R} \text{Posterior}(r|u,t) = 1 \tag{7}$$

Following normalization, the estimated reservation value, $\hat{r}$, is computed as the expected value over all hypotheses, weighted by their updated posterior probabilities:

$$\hat{r} = \sum_{r \in R} (r \cdot \text{Posterior}(r|u,t)) \tag{8}$$