BDI-based Opponent Modeling and Strategy Generation for Multi-Issue Negotiation (Student Abstract)

Tianzi Ma[1, 2], **Yulin Wu**[1, 2] *, **Hang Ren**[1, 2], **Xiaozhen Sun**[1, 2], **Shuhan Qi**[1, 2], **Xuan Wang**[1, 2] *
[1]Harbin Institute of Technology, Shenzhen, China
[2]Guangdong Provincial Key Laboratory of Novel Security Intelligence Technologies, Shenzhen, China
tianzima@stu.hit.edu.cn, yulinwu@cs.hitsz.edu.cn, 21b351008@stu.hit.edu.cn, 22b351013@stu.hit.edu.cn,
shuhanqi@cs.hitsz.edu.cn, wangxuan@cs.hitsz.edu.cn

Abstract

Accurately modeling opponent behaviors and integrating strategy are key challenges for multi-issue automated negotiation. Existing approaches often isolate preference learning or trend prediction and lack a unified cognitive structure with coordinated reasoning. This paper proposes a BDI (Belief-Desire-Intention)-based opponent modeling and strategy generation framework. The framework analyzes opponent responses (Belief), predicts preference weights and the utility function (Desire), and infers utilities of future offers (Intention). Building on this, we design a responsive strategy, enabling gradual concessions and balanced outcomes. Our main contributions are: D-MBUE in the Desire module, I-DABI in the Intention module, and the BDI Negotiator on top of the modeling modules. Experiments on 45 standard negotiation domains and against 12 representative opponents demonstrate the effectiveness of our BDI framework.

Code — https://github.com/13jqq/Supplementary_for_BDI_Negotiation.git

Introduction

Automated negotiation is a key approach for resource coordination and conflict resolution in multi-agent systems (Alshabi et al. 2007; Kraus 1997), aiming to reach mutually acceptable agreements through interaction under conflicting interests, and negotiation capabilities are becoming critical across domains. Negotiation often occurs under significant information asymmetry (Memon, Scoccia, and Autili 2025), so modeling the opponent's behavior and using it to guide offer generation and response are central challenges (Baarslag et al. 2016). However, most existing methods focus on isolated opponent-modeling tasks and do not integrate predictions into strategy generation, which weakens an agent's adaptability. We introduce a BDI-based opponent-modeling framework, *Belief* for state and behavior inference from offers, *Desire* for the opponent's utility function, and *Intention* for the next-offer direction, and a plug-and-play, behavior-aware negotiation strategy that uses these outputs to enhance performance. Through experiments we demonstrate that this framework achieves strong predictive accuracy and enhances negotiation outcomes.

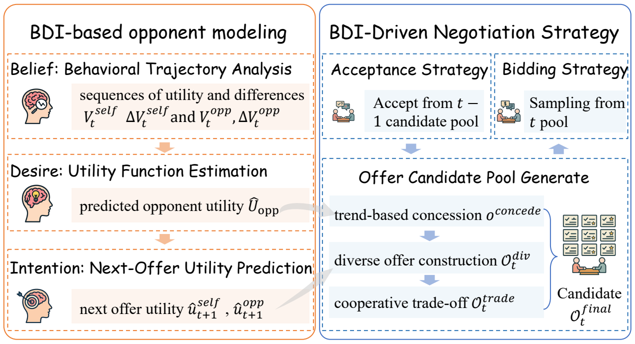

Figure 1: BDI-Based opponent modeling framework and BDI-Driven Progressive Negotiation Strategy

BDI-Based Opponent Modeling and Strategy

In multi-issue bilateral negotiation, two agents bargain over discrete issues $I = \{I_1, I_2, \ldots, I_n\}$ with option sets V_i. Each agent's utility function has a linear additive form: $U(\omega) = \sum_{i=1}^{n} w^i \cdot e^i(\omega_i)$, where w^i is the weight of issue i and $e^i(\cdot)$ maps options to evaluations. Negotiation follows the Alternating Offers Protocol (AOP) (Rubinstein 1982) under no prior information about the opponent.

We build a BDI-based opponent-modeling framework: Belief interprets negotiation history and both sides' offer trajectories; Desire (D-MBUE) predicts the opponent's utility function; Intention (I-DABI) estimates next-offer utility trends under dual perspectives; Strategy generates progressive, structurally decoupled, opponent-adaptive, concession-aware offers (see Figure 1).

Belief: Opponent Behavioral Trajectory Analysis

The Belief component captures the opponent's state understanding formed from the negotiation scenario and both sides' historical offers. Under information asymmetry, we use two perspectives: self (based on the agent's known utility function U_{self}) and opponent (based on the predicted utility function $\hat{U}_{\text{opp}}$ obtained from the Desire module), and, in each, build utility and first-difference sequences V_t^{self}, ΔV_t^{self} and V_t^{opp}, ΔV_t^{opp} to describe offer evolution.

---

*Corresponding authors.

Desire: Opponent Utility Function Estimation

The Desire module aims to predict the opponent's utility function $\hat{U}_{\text{opp}}(o) = \sum_{i=1}^{n} w_{\text{opp}}^{i} \cdot v_{\text{opp}}^{i}(o_i)$, where w_{opp}^{i} denotes the opponent's weight for issue i, and $v_{\text{opp}}^{i}(o_i)$ represents the opponent's evaluation of option o_i on that issue. We propose D-MBUE, a Bayesian utility function prediction method that constructs likelihood functions using multiple behavioral consistency assumptions and dynamically updates the posterior for real-time estimation. D-MBUE uses three behavioral-consistency likelihoods, concession monotonicity, early focus on high-utility regions, and similarity-based counteroffer improvement, and combines them with the prior to form a posterior weight $p_t \propto \mathcal{L}_1 \cdot \mathcal{L}_2 \cdot \mathcal{L}_3 \cdot P(\theta)$, which drives the update of the utility components and yields the predicted opponent utility function $\hat{U}_{\text{opp}}(o)$.

Intention: Opponent Next-Offer Utility Prediction

The Intention component models the opponent's potential offer utility for the next round, based on the joint bidding history and the previously predicted utility function. We propose I-DABI, a Bayesian predictor under dual perspectives (self utility and predicted opponent utility).

First, extract temporal features under U_{self} and $\hat{U}_{\text{opp}}$, and fit a Bayesian regressor to obtain predictive samples, yielding two sets $\left\{\hat{u}_{t+1}^{\text{self},(i)}\right\}_{i=1}^{n}$ and $\left\{\hat{u}_{t+1}^{\text{opp},(j)}\right\}_{j=1}^{n}$. Second, apply Dynamic Anchor Bayesian Regression (DABR), model utility with Bayesian linear regression $y_t = \boldsymbol{\beta}^{\top} X_t + \epsilon$ and fuse predictions with dynamic anchors $a_k = \text{mean}(y^{(k)})$, producing anchored predictive samples that stabilize trends. Third, cross-map the predictive samples to $\mathcal{O}$ under U_{self} and $\hat{U}_{\text{opp}}$ with tolerance ϵ: $\mathcal{O}_{\text{self}} = \left\{o \in \mathcal{O} : |U_{\text{self}}(o) - \hat{u}_{t+1}^{\text{self},(i)}| < \epsilon\right\}$, $\mathcal{O}_{\text{opp}} = \left\{o \in \mathcal{O} : |\hat{U}_{\text{opp}}(o) - \hat{u}_{t+1}^{\text{opp},(j)}| < \epsilon\right\}$, take the intersection $\mathcal{O}_{\text{final}} = \mathcal{O}_{\text{self}} \cap \mathcal{O}_{\text{opp}}$ and average utilities.

BDI-Driven Progressive Negotiation Strategy

We propose a BDI-driven progressive negotiation strategy that is structurally decoupled and concession-aware. It leverages the Desire module ($\widehat{U}_{\text{opp}}$) and the Intention module ($\hat{u}_{t+1}^{\text{self}}$) to guide offer generation and acceptance, while remaining plug-and-play with utility modeling modules.

For offer generation, we first produce a trend-aware concession offer o^{concede} (triggered by decreasing opponent utility under U_{self} or $\widehat{U}_{\text{opp}}$). We then construct the per-round candidate pool from three factors: previously proposed offers that meet the current threshold τ_t; the concession offer o^{concede}; and the predicted next offer $o_{t+1}^{\text{opp,pred}}$ from the Intention module if it also meets τ_t. A cooperative trade-off expands the pool with near-τ_t offers that achieve higher utility under $\widehat{U}_{\text{opp}}$, yielding $\mathcal{O}_t^{\text{final}}$. From this pool we sample offers according to $P(o) \propto \max\big(\varepsilon, \widehat{U}_{\text{opp}}(o)\big)$. For acceptance, the opponent's current offer is accepted if it lies in $\mathcal{O}_{t-1}^{\text{final}}$ and exceeds the threshold or reservation value.

Model	RMSE (Opp)	MAPE (Opp)	RMSE (Self)	MAPE (Self)
BDI	**0.1336**	**0.1151%**	**0.1013**	**0.2187%**
Concession	0.2042	0.2015%	0.1278	0.3061%
Monotonic	0.2028	0.2034%	0.1194	0.2715%
Stepwise	0.1647	0.1469%	0.1086	0.2678%

Table 1: Opponent Offer Utility Prediction Performance

Agent	UFun_mean	UFun_sem	ci_upper	ci_lower
BDI-Negotiator	**0.7530**	0.0163	0.7851	0.7203
Atlas3	0.6275	0.0141	0.6557	0.5996
Boulware	0.5737	0.0135	0.6002	0.5475
CUHKAgent	0.5188	0.0127	0.5437	0.4939
TheFawkes	0.5120	0.0173	0.5453	0.4784
MetaAgent2013	0.5117	0.0164	0.5433	0.4797
Conceder	0.4867	0.0107	0.5078	0.4660
TMFAgent	0.4407	0.0128	0.4662	0.4156
NaiveTitForTat	0.4119	0.0199	0.4501	0.3730
RandomDance	0.3609	0.0144	0.3883	0.3333
ParsAgent	0.3353	0.0262	0.3862	0.2844
PonPokoAgent	0.2712	0.0258	0.3232	0.2207
MiCRO	0.2371	0.0248	0.2859	0.1886

Table 2: Average self-utility against 12 opponents

Experimental Setup and Evaluation Metrics

We construct an experiment covering 45 multi-issue domains and run 2700 bilateral sessions; five agents (each with a different opponent-modeling technique but identical bidding/acceptance logic) negotiate against 12 state-of-the-art opponents in every domain, with 500 steps per session. We evaluate (i) opponent-modeling effectiveness, by predicting the utility of the opponent's next offer per round using RMSE and MAPE under both utility functions, (ii) negotiation performance, measured by average self-utility across all domains and opponents. As reported in Table 1, the BDI-based opponent modeling achieves the best scores across all four metrics, with RMSE/MAPE of 0.1336/0.1151% from the opponent's perspective and 0.1013/0.2187% from the self perspective. For negotiation performance, Table 2 shows that BDI-Negotiator attains the highest average self-utility (0.7530). These results validate the effectiveness of our integrated framework in terms of both prediction accuracy and negotiation performance.

Conclusion

This paper addresses the challenges where the opponent's utility function is unobservable and strategic behaviors are highly dynamic. We propose a BDI-based framework for opponent modeling and strategy generation. By introducing the BDI structure into negotiation scenarios, we construct a closed-loop opponent modeling framework that forms a cognitive chain from offer analysis to preference and utility function estimation, and further to offer trend prediction.

Acknowledgments

This research was funded in part by National Natural Science Foundation of China (General Program, No. 62376073), Shenzhen Natural Science Foundation (General Program, No. JCYJ20250604145644060), Colleges and Universities Stable Support Project of Shenzhen (No. GXWD20220811170225001), and Guangdong Provincial Key Laboratory of Novel Security Intelligence Technologies (No. 2022B1212010005).

References

Alshabi, W.; Ramaswamy, S.; Itmi, M.; and Abdulrab, H. 2007. Coordination, cooperation and conflict resolution in multi-agent systems. In *Innovations and advanced techniques in computer and information sciences and engineering*, 495–500. Springer.

Baarslag, T.; Hendrikx, M. J.; Hindriks, K. V.; and Jonker, C. M. 2016. Learning about the opponent in automated bilateral negotiation: a comprehensive survey of opponent modeling techniques. *Autonomous Agents and Multi-Agent Systems*, 30: 849–898.

Kraus, S. 1997. Negotiation and cooperation in multi-agent environments. *Artificial intelligence*, 94(1-2): 79–97.

Memon, M. A.; Scoccia, G. L.; and Autili, M. 2025. A systematic mapping study on automated negotiation for autonomous intelligent systems. *Automated Software Engineering*, 32(2): 1–46.

Rubinstein, A. 1982. Perfect equilibrium in a bargaining model. *Econometrica: Journal of the Econometric Society*, 97–109.