Strategic Tool Enhanced AI Agent for Multi-Issue Negotiation (Student Abstract)

Daiki Kitashima[1, 2], Ryota Higa[2, 3], Katsuhide Fujita[1, 2]

[1]Tokyo University of Agriculture and Technology, 2-24-16 Koganei, Tokyo, 184-8588, Japan,
[2]National Institute of Advanced Industrial Science and Technology,
[3]NEC Corporation
kitashima@katfuji.lab.tuat.ac.jp, r-higaryouta@nec.com, katfuji@cc.tuat.ac.jp

Abstract

Automated negotiation, a form of interaction among autonomous agents, plays a central role in multi-agent systems, yet the application of large language model (LLM) in this domain remains underexplored. An LLM can serve as a meta-strategist, adaptively selecting explicit strategies for execution by external strategic tools based on its capabilities. We propose a negotiation AI agent equipped with explicit strategic tools, including time-dependent and tit-for-tat negotiation strategies. Our results show that strategic tool enhanced negotiators achieve approximately 16% higher average utility compared with baseline, latest LLM negotiators.

Introduction

Automated negotiation is essential in multi-agent systems, as it enables resource allocation and conflict resolution. Existing approaches primarily focus on rule-based, optimization, and reinforcement learning methods (Higa et al. 2023). The application of large language model (LLM), with their reasoning and adaptive capabilities, to automated negotiation remains underexplored (Bianchi et al. 2024).

Although LLMs excel at generating fluent dialogue, they often struggle to adapt strategic behavior to opponents in one-shot, multi-issue negotiations. We therefore consider the LLM as a meta-strategist that adaptively calls multiple explicit strategies executed by an external strategic tool.

We propose a negotiation AI agent that leverages external strategic tools to provide strategy-based negotiation actions. By integrating these tools, the LLM can maintain coherent negotiation strategies while retaining flexibility in adaptive, context-sensitive communication. In this hybrid design, the LLM generates dialogue, while the strategic tools guide concession patterns, resulting in more coherent and effective negotiations. Our findings show that incorporating explicit strategic tools yields better outcomes than LLM negotiators that do not use external strategic tools.

Bilateral Multi-Issue Negotiation

We consider a bilateral negotiation setting between two agents. The negotiation domain is defined by a set of n issues, $I = \{I_1, \ldots, I_n\}$, where each issue I_i has k_i possible values, $V_i = \{v_1^i, v_2^i, \ldots, v_{k_i}^i\}$. During negotiation,

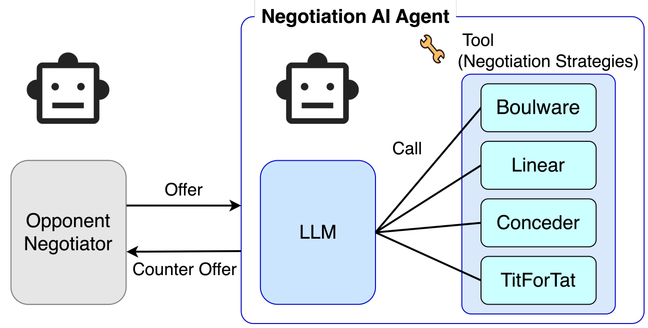

Figure 1: Overview of the proposed negotiation architecture. The proposed negotiator can invoke different negotiation strategies as external strategic tools.

each agent proposes a bid, selecting one value for each issue, represented as $\omega = (v_{c_1}^1, v_{c_2}^2, \ldots, v_{c_n}^n)$, where $v_{c_i}^i \in V_i$. The set of all possible bids is denoted by Ω, with size $|\Omega|$. Each agent has a private preference structure defined by a utility function $U(\cdot)$, where the utility of a bid ω is defined as a weighted sum of evaluations over issues: $U(\omega) = \sum_{i=1}^{n} w_i \cdot eval(v_{c_i}^i)$ Here, $eval(v_{c_i}^i)$ denotes the local evaluation value assigned by an agent to option $v_{c_i}^i$, which is scaled to lie within the range $[0, 1]$. The parameter $w_i \geq 0$ represents the weight of the issue I_i, with $\sum_{i=1}^{n} w_i = 1$. Thus, the utility values are bounded as $U(\omega) \in [0, 1]$.

Negotiation follows an alternating-offers protocol (Higa et al. 2023). At each turn, an agent selects one of the following actions: *Accept* (accepting the most recent bid of the opponent, which immediately concludes the negotiation with agreement), *Offer* (proposing a bid $\omega \in \Omega$ to the opponent), or *Terminate* (ending the negotiation without agreement). The process continues until an agreement is reached through acceptance, explicit termination occurs, or a predefined deadline is met.

Strategic Tool Enhanced Negotiation Strategy with Large Language Models

As shown in Figure 1, the negotiation AI agent interacts with the negotiation environment while consulting a strategic tool that provides strategy-based actions. We developed a single

Negotiation Agent	Boulware	Linear	Conceder	TitForTat1	TitForTat2	AgentK	HardHeaded	Atlas3	AgentGG	Average
o4-mini (w/o strategic tool)	0.336	0.623	0.857	0.884	0.744	**0.356**	0.187	0.515	0.376	0.542
o4-mini (w/ strategic tool)	**0.396**	**0.628**	**0.879**	**0.938**	**0.942**	0.346	0.233	**0.661**	**0.628**	**0.628**
GPT-4.1 mini (w/ strategic tool)	0.240	0.462	0.764	0.640	0.552	0.242	**0.259**	0.311	0.332	0.422

Table 1: Negotiation results between evaluated negotiators (rows) and their opponents (columns) in EnergySmall_A, a domain with a domain size of 15,625 and a conflict level of 0.435. Each entry reports the mean utility achieved by the evaluated negotiator. Average denotes the mean utility over all negotiations with all opponents.

strategic tool that can be invoked by the LLM during negotiation, which implements multiple predefined strategies and returns either an offer or a response based on the current negotiation state.

The LLM prompt included negotiation-related information (e.g., the negotiation domain, utility functions, and the opponent's proposal) along with a description of the strategic tool. The LLM was instructed to use the strategic tool at least once per turn and to output the next action (acceptance or counteroffer) in an XML-like format. This design ensures that the LLM incorporates the tool's strategic output into its reasoning when determining the next action. To select a strategy, the LLM specifies an argument when invoking the strategic tool, allowing flexible switching among strategies within a single turn. If the LLM generated an action without invoking the strategic tool, the output process was repeated until a tool call was made. This ensured that all actions were grounded in the results of the strategic tool.

Five strategies were implemented, specified through arguments and grouped into two categories: **Time-dependent** (adjusting concession rates as negotiations progress, including Boulware, linear, and conceder) and **Tit-for-Tat** (adapting to the opponent's previous offers, with variants considering either one- or two-step history).

We employed NegMas[1] as the negotiation platform and implemented our negotiators utilizing the Google Agent Development Kit (Google ADK)[2] to support agent development and integration.

Evaluations

Experimental Setup

The negotiation domain used in our experiments is EnergySmall_A. This domain consists of six issues, each with five discrete options, yielding a total of $|\Omega| = 5^6 = 15,625$ possible offers, with a domain conflict level of 0.435. It provides sufficient complexity to evaluate the strategic consistency and adaptability of the agents. Negotiations followed an alternating-offers protocol, with a maximum of 80 turns per agent. In each turn, agents either responded to the opponent's offer or made a new proposal. In all experiments, the LLM-based negotiator was the first mover. Each experiment was repeated 10 times to ensure robustness.

Baseline We compare our method with the latest LLM-based negotiators that do not employ any strategic tools. This baseline enables us to compare and verify the impact of strategic tools on negotiation performance.

[1]https://github.com/yasserfarouk/negmas
[2]https://google.github.io/adk-docs/

Experimental Result

We selected o4-mini as the LLM with the best performance after evaluating it against other LLMs (GPT-4.1 nano and GPT-4.1 mini) based on the achieved utility values.

The results are summarized in Table 1. On average, the negotiation AI agent with strategic tools achieved a higher utility (0.628) than the latest negotiation AI agent (0.542), representing a relative improvement of approximately 16%. In addition, our negotiator achieved a 100% agreement rate against all opponents. With an average utility of 0.560, these opponents outperformed the latest LLM, whereas our negotiator outperformed the opponents.

The increase in utility values is attributed to the suppression of premature concessions and excessively stubborn bargaining. No substantial differences were confirmed in the average number of turns to conclude negotiations between the strategic tool augmented and the latest LLMs. However, for AgentGG, where the utility improved substantially, the mean number of turns changed only slightly ($48.1 \pm 24.8 \rightarrow 46.9 \pm 15.6$), with the most notable effect being a marked reduction in variance. Similar trends were observed for all negotiators without the strategic tool whose standard deviation exceeded 10, indicating that declines in utility from either conceding too early or refusing to concede were effectively mitigated. Furthermore, except for the Conceder strategy, the SD of utility values decreased, suggesting that the strategic tool mitigates the inherent randomness of the LLM and enables more strategy-aware negotiations.

Beyond accurate concession control, the ability to select from multiple strategies further enhances the negotiator's adaptability. By choosing strategies based on the opponent's behavior and negotiation context, the LLM can dynamically adjust its concession patterns, improving stability and overall performance.

Conclusions

We proposed a negotiation AI agent equipped with explicit strategic tools, including time-dependent and tit-for-tat strategies. Our results demonstrate that strategic tool augmented negotiators achieve higher average utility than the latest LLMs. Future work will explore richer negotiation domains and investigate dynamic selection among multiple strategic tools to further enhance adaptability.

References

Bianchi, F.; Chia, P. J.; Yuksekgonul, M.; Tagliabue, J.; Jurafsky, D.; and Zou, J. 2024. How well can LLMs negotiate? NEGOTIATIONARENA platform and analysis. In *Proceedings of ICML'24*.

Higa, R.; Fujita, K.; Takahashi, T.; Shimizu, T.; and Nakadai, S. 2023. Reward-Based Negotiating Agent Strategies. In *Proceedings of AAAI 2023 Conference*, 11569–11577.