EOHNegotiator: An Adaptive Agent for Sequential Multi-Deal Negotiations

EOH Team
Eliran Eiluz, Hallel Weinberg, Oriya Cohen

June 9, 2025

Abstract

The Automated Negotiation League (ANL) 2025 challenge focuses on sequential multi-deal negotiation, where an agent must negotiate with multiple opponents in sequence to achieve an optimal combined outcome. In response to this challenge, we present EOHNegotiator, an adaptive agent designed to navigate the complexities of this one-to-many negotiation environment. The agent's architecture is built on three core principles: a clear distinction between strategies for the center and edge roles, a dynamic, time-based concession approach for bidding and acceptance, and an adaptive mechanism that adjusts key parameters based on historical negotiation performance. This report details the design of EOHNegotiator, outlining its coordination, bidding, and acceptance strategies, and concludes with lessons learned from its development.

1 Introduction

The Automated Negotiating Agent Competition (ANAC) is an international tournament designed to foster research in automated negotiation by presenting a new and complex challenge each year. The 2025 Automated Negotiation League (ANL) focuses on the domain of sequential multi-deal negotiation. In this setting, a "center agent" must negotiate one-on-one with multiple "edge agents" sequentially, where its final utility is determined by the combination of all deals secured. This creates a complex strategic problem requiring agents to look ahead, manage risk, and coordinate outcomes across multiple negotiations, mirroring real-world scenarios like supply chain management and procurement.

The agent must be general enough to handle diverse scenarios, from coordinating social events to acquiring a target quantity of goods. Success requires more than just optimizing a single deal; it demands a holistic strategy that balances the potential rewards of future negotiations against the certainty of present ones. Our agent, EOHNegotiator, is designed to address this challenge through an adaptive, role-aware framework.

2 The Design of EOHNegotiator

The design of EOHNegotiator is centered around three primary components as recommended by the competition guidelines: a coordination strategy for the center agent role, a bidding strategy, and an acceptance strategy. The agent's logic is fundamentally adaptive, dynamically adjusting its behavior based on its role, the passage of time within a subnegotiation, its progress through the overall sequence of negotiations, and its historical success rate.

2.1 Coordination Strategy (Center Agent)

The core challenge for the center agent is to manage the interdependencies between sequential negotiations. EOHNegotiator's coordination strategy is primarily handled by the `_update_strategy` method, which is triggered at the conclusion of each subnegotiation.

The central mechanism for coordination is the dynamic recalculation of the agent's long-term goal. After an agreement is (or is not) reached, the agent calls the helper function `find_best_bid_in_outcomespace`. This function performs a crucial task:

1. It considers all deals from past negotiations as "locked in" or fixed.
2. It then re-evaluates all possible outcomes for the *remaining* future negotiations.
3. It identifies the new single best combined outcome that is still achievable. This outcome becomes the new `target_bid`.

This ensures the agent is never working towards an obsolete goal and always re-optimizes its strategy based on the current state of the world.

Furthermore, the agent's proposal strategy as a center agent (`_propose_as_center`) incorporates its position in the sequence. In the first half of the negotiations, it intentionally reduces its concession factor by 30% to be more aggressive and anchor the negotiations around high-utility outcomes. In the latter half, it increases its concession factor by 20%, recognizing the need for greater flexibility as its options dwindle.

2.2 Bidding Strategy

The agent's bidding strategy is driven by a time-based concession model. For both center and edge roles, the agent calculates a `concession_factor` that determines how far below its optimal `target_utility` it is willing to offer a bid.

The concession is non-linear, governed by the formula:

$$\text{concession} \propto (\text{time\_pressure})^2 \times \text{concession\_rate}$$

This ensures the agent concedes very little at the beginning of a negotiation, holding firm on its high-utility goals, and concedes much more rapidly as the deadline approaches.

The bidding strategy is also role-dependent:

- **As Center Agent:** It behaves more aggressively, moderated by the sequential progress heuristic described previously.
- **As Edge Agent:** It is designed to be more cooperative, multiplying its concession factor by a `self.edge_cooperation` parameter, making it more willing to yield to find an agreement.

Once an `acceptable_utility` is determined, the agent uses the `_find_bid_with_utility_from_current_space` method to find a concrete bid to propose. This method searches the outcome space of the current subnegotiation for a bid that closely matches the agent's current utility target.

2.3 Acceptance Strategy

The decision to accept or reject an offer is governed by a dynamic `acceptance_threshold`. This threshold starts high (close to the agent's target utility) and decreases over time, mirroring the concession logic of the bidding strategy.

A key feature of EOHNegotiator is its adaptive threshold mechanism, which modifies its acceptance behavior based on past performance:

- If the agent's overall `success_rate` drops below 30%, it becomes more lenient, lowering its acceptance threshold by 20% to increase the chance of securing a deal.
- Conversely, if its `success_rate` exceeds 70%, it becomes more demanding, raising its threshold by 10% to capitalize on its strong performance.

This creates a feedback loop where the agent learns from its history across a single multi-deal session.

Additionally, when acting as an **edge agent**, it employs a simple form of opponent modeling. It checks if the utility of the opponent's last two offers has been increasing. If the opponent is actively conceding, EOHNegotiator becomes slightly tougher, raising its own acceptance threshold to see if a better offer is forthcoming.

3 Evaluation

The agent was developed and tested locally using the provided `run_a_tournament` helper script against the standard set of ANL 2025 agents (e.g., `Boulware2025`, `Conceder2025`). The primary goal of testing was to ensure robustness and logical correctness across different scenarios inspired by the competition, such as the Dinner, Target Quantity, and Job Hunt scenarios.

A significant challenge identified during development is the computational complexity of the center agent's coordination strategy. The reliance on `all_possible_bids_with_agreements_fixed`, which performs a brute-force search of the remaining outcome space, is computationally expensive. In scenarios with large outcome spaces, this could potentially lead to the agent timing out before it can calculate its optimal strategy. This represents the agent's most significant weakness.

Qualitatively, the agent performs well against predictable, time-based agents due to its own similar but adaptive logic. Its performance against highly irrational (`Random2025`) or extremely stubborn (`Boulware2025`) agents is more mixed, as its simple opponent modeling is not sufficient to fully exploit their behavior. The final score is calculated as an average of its performance in both center and edge roles, as specified in the competition rules.

4 Lessons and Suggestions

The development of EOHNegotiator provided several key insights into the challenges of sequential multi-deal negotiation.

1. **The Centrality of Re-planning:** The most critical function for the center agent is its ability to dynamically re-calculate its optimal path after every subnegotiation. A static plan is guaranteed to fail.
2. **The Computational Barrier:** The primary obstacle to optimal planning is the "combinatorial explosion" of possible future outcomes. The agent's current brute-force approach highlights this limitation. A clear path for future improvement would be to replace this with a more efficient sampling-based heuristic, such as evaluating a limited, random sample of future possibilities.
3. **Adaptation is Key:** A simple, fixed strategy is easily exploited. The agent's performance improved significantly after implementing the adaptive acceptance threshold based on success rate, demonstrating the value of learning from past performance within a single multi-deal session.

Conclusions

EOHNegotiator is a robust, adaptive agent designed for the ANL 2025 competition. Its core strength lies in its adaptive time-based concession strategy and its dynamic re-planning mechanism for the center agent role. While its strategic planning can be computationally intensive, it provides a solid foundation and a clear demonstration of the core principles required to succeed in a sequential multi-deal negotiation environment.