SmartNegotiator: A Boulware-Inspired Adaptive Agent for Multi-Party Negotiation

1 Introduction

We present *SmartNegotiator*, an automated negotiation agent designed for the ANL2025 competition. Our agent combines a Boulware-style concession strategy with opponent modeling and sequential coordination strategies to adapt its behavior dynamically based on negotiation progress and opponent concession patterns. The agent functions as both a center and edge agent in multilateral negotiations and adapts its utility thresholds accordingly. SmartNegotiator's architecture is designed to balance long-term global objectives with short-term negotiation tactics, making it highly adaptable to rapidly changing negotiation conditions. The design prioritizes coordination, adaptability, and robustness in high-stakes negotiation environments.

2 Coordination Strategy

In multi-party settings, a center agent must coordinate multiple sequential bilateral negotiations. SmartNegotiator addresses this challenge by maintaining a shared state that includes an index for the current subnegotiation, a dictionary mapping subnegotiation IDs to their status, and an updated target bid.

The agent computes a global target bid by optimizing the utility over the space of all feasible multi-issue outcomes, conditioned on previously fixed agreements. This bid represents an optimal tuple across all subnegotiations and guides both proposals and acceptances in the current and future negotiations.

When evaluating a potential bid in the current subnegotiation, the agent constructs the Cartesian product of:

- Fixed agreements from past negotiations
- The current offer
- The set of possible values for remaining negotiations

This ensures compatibility and strategic alignment across the negotiation sequence. The coordination logic thus enables the agent to locally optimize each subnegotiation while adhering to a globally consistent strategy.

Additionally, the agent incorporates a Boulware-style threshold that is time-dependent. This time-adaptive behavior introduces a softening effect on the coordination strategy as negotiations progress, ensuring the agent remains assertive early but making the agent more acceptive towards the end of the negotiation timeline.

3 Bidding Strategy

SmartNegotiator's bidding strategy integrates temporal concession patterns and opponent modeling to construct an adaptive utility threshold each turn. The strategy blends a **local threshold**, based on the number of steps in the current negotiation, and a **global threshold**, derived from the index of the current subnegotiation.

Formally, the global threshold follows a Boulware-inspired formula:

where t, the normalized negotiation index, and hyperparameters β, k control the steepness and floor of the threshold.

The agent also employs opponent modeling by estimating the opponent's concession behavior using an exponentially weighted moving average (EWMA) over changes in the utility of received offers. This estimation adjusts the threshold up or down, bounded within ± 0.2.

All bids in the outcome space are filtered to retain only those whose utility exceeds the adjusted threshold. The agent then samples a bid probabilistically, with the selection likelihood proportional to its utility. This stochastic component avoids local optima and fosters exploration, especially useful in uncertain environments.

In the beginning of each negotiation, the agent will propose the best possible bid. If the best available bid is None, indicating that reaching any agreement is no longer beneficial, the agent will end the negotiation instead of proposing a new offer. This ensures that the agent does not engage in unnecessary or unprofitable negotiations.

4 Acceptance Strategy

The acceptance mechanism mirrors the dynamic and adaptive structure of the bidding logic. Upon receiving an offer, SmartNegotiator computes its utility by evaluating it within the context of the fixed agreements and compatible values for remaining negotiations.

The acceptance threshold is constructed similarly to the bidding threshold, combining both local and global dynamics and adapting based on the opponent's observed concession behavior. An offer is accepted if its utility exceeds the product of this threshold and the utility of the agent's target bid for the current subnegotiation.

This ensures that the agent does not accept offers that would compromise its overall strategic position while remaining flexible against concessionary opponents. If no threshold is met, the offer is rejected. This strict yet adaptive mechanism ensures robustness against both exploitative and overly rigid strategies.

5 Implementation and Parameters

The agent was implemented in Python, using the NegMAS negotiation platform. Several hyperparameters were tuned empirically:

- $\beta = 2.68$: Controls the steepness of the Boulware-style global threshold.
- $k = 0.7$: Sets the lower bound of concession.
- min acceptable utility $= 0.61$: Floor value for the threshold.
- decay $= 0.03$: Step-based decay for the local threshold.
- $\alpha = 0.5$: Smoothing factor for EWMA-based opponent modeling.

The architecture supports both edge agents (who negotiate only once) and center agents (who negotiate sequentially across all issues). The agent's internal memory tracks past offers, opponent behavior, and utility changes, allowing it to detect patterns and refine its strategy as the tournament progresses.

6 Conclusion

SmartNegotiator demonstrates a principled yet flexible approach to automated negotiation in complex multi-agent environments. By integrating Boulware-style concession with real-time opponent modeling, coordinated subnegotiation management and probabilistic bid selection, the agent balances local negotiation efficacy with global coordination goals.

This multi-layered strategy makes SmartNegotiator robust to a range of opponent behaviors and negotiation topologies. Future improvements may explore predictive modeling of opponent types, multi-issue trade-off heuristics, and reinforcement learning to further optimize multi-turn decision-making.