

OverlapConceder

Hossein Savari, University of Tehran, h.savari@ut.ac.ir

June 22, 2026

Abstract

In this report we describe *OverlapConceder*, an automated negotiating agent designed for the ANAC 2026 ANL league. The agent combines a time-based concession strategy with bounded target utility, a dual-criterion acceptance policy, and a lightweight opponent model based on weighted offer similarity. All agent parameters were tuned using Bayesian optimisation against a set of opponents and scenarios. The final configuration achieves a robust balance between individual gain and concealment.

1 Introduction

This report explains the *OverlapConceder* agent, a negotiating agent that combines a dynamic concession model, opponent-aware bidding, and an adaptive acceptance rule.

The agent does so by maintaining a time-varying *target utility range* from which its offers are generated, which allows the agent to concede slowly while pursuing high gains when an opponent model is available. All continuous parameters were optimized using Gaussian-process Bayesian optimization over a test suite. The resulting agent balances concession, exploitation of estimated opponent preferences, and the concealment of its own preferences.

2 The Design of OverlapConceder

Figure 1 shows the high-level decision cycle. On each turn, the agent first updates its opponent model using the received offer, then decides whether to accept, and if not, it generates a new bid using its concealing bidding strategy.

The agent’s behaviour is decided by nine tunable parameters that are listed in Table 1. They control the concession speed (β_a, β_b), the target utility scaling, the acceptance threshold, the opponent-model weighting (γ), the utility mapping, and the number of candidate bids considered.

3 Concealing Bidding Strategy

The concealing bidding strategy generates offers that lie inside a time-varying target utility interval. Let $t \in [0, 1]$ be the relative negotiation time. The interval bounds are

$$U_{\text{low}}(t) = \frac{u_{\text{max}} - \Delta \cdot t^{\beta_a}}{s} + c, \quad (1)$$

$$U_{\text{high}}(t) = \frac{u_{\text{max}} - \Delta \cdot t^{\beta_b}}{s} + c + b, \quad (2)$$




Figure 1: Decision cycle of the OverlapConceder agent.

Table 1: Agent parameters and their final optimised values.

Parameter	Description	Value
β_a	Lower time exponent (concession speed)	1.2344
β_b	Upper time exponent	1.2922
s	Target scale factor	0.8211
c	Target additive shift	0.9174
b	Upper limit shift	0.9174
α	Advantage over opponent estimate	0.3871
γ	Decay factor for opponent model	0.9869
$u_{\min}$	Minimum opponent utility guess	0.4181
$u_{\max}$	Maximum opponent utility guess	0.9911
k	Number of top welfare bids	-

where $u_{\max}$ is the agent’s maximum achievable utility, $\Delta = u_{\max} - u_{\text{res}}$ is the surplus above the reservation value, and s, c are scale and shift constants. Because $\beta_a < \beta_b$, the interval shrinks as time advances, forcing more concession (Figure 2).




Figure 2: Concession curves with $\beta_a = 1.5$, $\beta_b = 2$, $s = 1.5$, $c = 0.1$, $b = 0.05$. The two curves diverge for $0 < t < 1$ because $\beta_a < \beta_b$, creating a target window that the agent samples from. The agent draws a random target $U_{\text{target}} \sim \mathcal{U}(U_{\text{low}}, U_{\text{high}})$ and bids an outcome with utility as close as possible.

A random target U_{target} is drawn uniformly from the current interval. If the opponent model is available, the agent selects among all rational outcomes whose utility falls inside $[U_{\text{low}}, U_{\text{high}}]$ the one that maximises the *combined welfare* $u_{\text{self}} \cdot \hat{u}_{\text{opp}}$. The top- k outcomes are considered, and then the bid is chosen from this set to add randomness. For this competition however this part has been disabled and it only uses pure concession.

If no opponent model exists, then the agent falls back to pure concession, it scans the list of rational outcomes (sorted by decreasing own utility) and picks the first one with utility at least U_{target} . In both cases the bid doesn’t reveal the exact target and remains relatively unpredictable.

4 Acceptance Strategy

The acceptance strategy uses two criteria. An opponent’s offer is accepted if

1. $u_{\text{self}}(\text{offer}) \geq U_{\text{low}}(t)$, the offer meets the lower concession bound, or
2. $u_{\text{self}}(\text{offer}) > \hat{u}_{\text{opp}}(\text{offer}) + \alpha$, where $\hat{u}_{\text{opp}}$ is the estimated opponent utility and α is an advantage threshold.

The first condition ensures that the agent does not miss offers that are already within its acceptable range. The second condition allows acceptance of an offer that is below the lower bound but is sufficiently better for the agent than for the opponent and it prevents the agent from rejecting a deal that is clearly in its own favour. The parameter α controls how much better the agent’s utility must be before an acceptance in the second case.

5 Opponent Model

The opponent model estimates the opponent’s utility $\hat{u}_{\text{opp}}(x)$ for any outcome x using the history of received offers. For an outcome with n issues,

$$\hat{u}_{\text{opp}}(x) = u_{\min} + (u_{\max} - u_{\min}) \cdot s(x),$$

where $s(x)$ is a weighted average similarity:

$$s(x) = \frac{\sum_{i=0}^{m-1} w_i \cdot \text{sim}(x, \mathbf{o}_{m-1-i})}{\sum_{i=0}^{m-1} w_i}, \quad w_i = \gamma^i,$$

$\mathbf{o}_j$ is the j -th opponent offer (oldest first), and $\text{sim}(x, \mathbf{o}) = \frac{1}{n} \sum_{j=1}^n \mathbf{1}[x_j = o_j]$ is the fraction of issues on which the two outcomes agree. The decay factor $\gamma \in (0, 1]$ gives more weight to recent offers, allowing the model to track a possibly non-stationary opponent. The constants $u_{\min}$ and $u_{\max}$ map the similarity score onto a plausible utility range which prevents extreme extrapolation.

However, the model has some limitations, it assumes issue-level independence and cannot consider issue interactions or trade-offs. In scenarios with issue dependencies, the model may produce inaccurate predictions, future work could address this.

6 Evaluation

The nine parameters were tuned using Bayesian optimisation with Gaussian processes (scikit-optimize). The objective function evaluated each parameter configuration on a fixed set of 96 random test cases. Each test case paired the agent against one of four opponent types (Boulware, BOA, MAP, Simple) on a scenario drawn from a set of negotiation domains, with a random number of steps (5–50) and a random order of play.

We used 60 initial random points and 60 further iterations of the GP-minimise algorithm with expected-improvement acquisition, a noise level of 0.05, and a fixed random seed for reproducibility. The best parameter set obtained is shown in Table 1.

7 Lessons and Suggestions

The development of OverlapConceder showed us a few insights:

- **Time-based concession with a target range** is more robust than a single target curve, as it gives the agent the ability to show random behaviour and exploit opponent information while still making progressive concessions.
- **The opponent model**, provides sufficient signal to distinguish between purely selfish bids and mutually beneficial ones.
- **The dual-criterion acceptance rule** prevents deadlocks when the opponent makes a generous offer that falls slightly below the agent’s lower bound, and it avoids over-concession by requiring a clear advantage when accepting below the bound.
- **Bayesian optimisation** is an effective tool for tuning negotiation strategies, but as the objective function noisy, it is better to add a moderate noise level for stabilised training.
- **Suggestions** we can replace the simple similarity model with a learned utility function while retaining the welfare-bidding method, and could adapt the concession exponents online based on the opponent’s behaviour.

Conclusions

We presented the *OverlapConceder*, a negotiating agent that combined a concealed time-dependent concession strategy, an opponent-aware bidding mechanism, and a lightweight opponent model. The agent’s parameters were tuned using Bayesian optimisation. The agent’s final parameters should balance concession, the exploitation of the opponent’s estimated preferences, and the concealment of its own preferences.