

LionelNeg: A Frequency–Modelling Boulware Agent for the ANAC 2026 Automated Negotiation League

Liu Wei

Nanjing University of Science and Technology
w79370103ei@gmail.com

June 2026

Abstract

We describe **LionelNeg**, our entry to the ANAC 2026 Automated Negotiation League (ANL). The 2026 league scores each negotiator by the sum of two terms, *Advantage* (the utility secured above the reservation value) and *Concealing* (a normalised, rank-correlation measure that simultaneously rewards modelling the opponent well and being hard to model). LionelNeg couples a disciplined Boulware time-based concession curve with a multi-criteria offer selector that blends self-utility, an opponent-frequency estimate and a Nash-product cooperation term, a monotone ACNEXT-style acceptance rule with a downward-ratcheted threshold, and a combined Smith + distribution frequency opponent model that is also exposed for the Concealing metric. In head-to-head tournaments over the seven official scenarios, LionelNeg ranks first ahead of BOANEG and MAPNEG, though by a narrow margin ($\Delta \approx 0.005$ in total score). We also report a clear negative result: across five independent design attempts and an LLM-driven evolutionary search, Concealing proved to be structurally capped, so all reliable score gains came from the Advantage side.

1 Introduction

ANL 2026 is a bilateral, alternating-offers negotiation over a discrete, multi-issue outcome space. Unlike the 2025 multi-deal league, each negotiation is a single session, but the scoring function changes the objective fundamentally. For a negotiator i the score is

$$\text{Score}_i = \underbrace{u_i(\omega^*) - r_i}_{\text{Advantage}} + \underbrace{\frac{a_i}{a_i + a_j}}_{\text{Concealing}}, \quad a_k = \frac{1 + \tau(u_{\bar{k}}, \hat{u}_k)}{2}, \quad (1)$$

where ω^* is the agreement, r_i the reservation value, $\tau(\cdot, \cdot)$ is Kendall’s rank correlation over the outcome space, $u_{\bar{k}}$ is the *true* utility function of k ’s opponent, and $\hat{u}_k$ is the model that agent k *exposes* of its opponent. The Concealing term is therefore a zero-sum ratio between the two parties: it rewards an agent both for modelling its opponent accurately (the numerator a_i) and for making its own preferences hard to recover (which lowers the opponent’s accuracy a_j).

This two-objective structure creates a tension that drives our design. Securing Advantage requires repeatedly proposing high-utility outcomes, which inevitably reveals our preference ordering and thus *helps* the opponent’s model of us. LionelNeg addresses the tension pragmatically: it pursues Advantage with a principled concession schedule and opponent-aware offer selection, while keeping an accurate opponent model for the Concealing numerator and varying its own offers as a design-level attempt to blunt the denominator. Section 7 reports our finding that the Concealing term is, in practice, near a structural ceiling, which reframes the whole competition as an Advantage-maximisation problem under a soft information-leakage penalty.

2 The Design of LionelNeg

LionelNeg follows the classic BOA decomposition (Bidding, Acceptance, Opponent model) with a thin domain-preparation layer. On the first preference-change callback it caches the full outcome space: it precomputes and stores every outcome utility, the achievable range $[u_{\min}, u_{\max}]$ and the reservation value r , and sorts outcomes by utility for fast inverse lookups (capped at a configurable pool size for very large domains such as CAR, which has 1.5×10^4 outcomes). It also flags *extreme scenarios* — domains such as NICEORDIE where more than 85% of outcomes lie within 0.10ρ of the reservation value ($\rho = u_{\max} - r$) — and switches to a more conceding floor in those cases so that the single win-win outcome remains reachable. All subsequent decisions reuse this cache, keeping per-step cost low.

3 Concealing Bidding Strategy

Target utility. Our aspiration follows a Boulware polynomial of relative time $t \in [0, 1]$,

$$\text{asp}(t) = 1 - t^5, \quad u^{\text{tgt}}(t) = f + \text{asp}(t) (u_{\max} - f), \quad f = r + 0.10\rho, \quad (2)$$

so the agent holds out near $u_{\max}$ for roughly the first two thirds of the session and then concedes quickly. In extreme scenarios the floor f drops to r .

Candidate selection. Rather than offer the single outcome closest to u^{tgt} , LionelNeg builds a candidate set within a time-shrinking band $[u^{\text{tgt}} - \beta, u^{\text{tgt}} + \beta]$, augments it with recently received offers above a guarded floor, and scores each candidate ω by

$$\begin{aligned} S(\omega) = & w_s \hat{s}(\omega) + w_o \hat{o}(\omega) + w_p \hat{s}(\omega) \hat{o}(\omega) + 0.20 \text{close}(\omega) \\ & + 0.07 \text{div}(\omega) + w_n \text{nov}(\omega) + w_e \text{exact}(\omega) - 0.14 \text{rep}(\omega), \end{aligned} \quad (3)$$

where $\hat{s}$ is our normalised utility, $\hat{o}$ the opponent estimate (Section 5), and $\hat{s} \hat{o}$ is a Nash/Pareto cooperation term that favours mutually beneficial outcomes. The weights are time-varying: $w_s = 0.45 - 0.12t$ decreases while $w_o = 0.25 + 0.35t$ increases, so the agent is greedy early and cooperative late, mirroring the concession schedule. The *novelty* and *diversity* terms (nov, div) reward offers that vary the values used across issues, avoiding sterile repetition and maintaining variety in the observable offer sequence. A hard floor forbids proposing below $u^{\text{tgt}} - 0.20\rho$, which prevents the opponent-aware terms from dragging us onto the opponent’s favourite outcome too early.

4 Acceptance Strategy

LionelNeg uses an ACNEXT-style rule: it accepts the incoming offer ω when its utility meets a threshold defined by our own next planned move,

$$\theta(t) = \max(f_0, u(\omega^{\text{next}}), 0.87 u^{\text{tgt}}(t)), \quad \text{accept iff } u(\omega) \geq \theta(t), \quad (4)$$

with a small floor $f_0 = r + 0.01\rho$ (raised to $r + 0.15\rho$ in extreme scenarios). The $0.87 u^{\text{tgt}}$ term keeps the bar tied to our aspiration so the opponent-aware bidding terms cannot quietly lower it into a premature deal. Crucially, we apply a *monotone downward ratchet* (A1): the realised threshold is clamped to the running minimum of all thresholds seen, $\theta \leftarrow \min(\theta, \theta_{\min})$, and $\theta_{\min}$ is updated thereafter. Because the Boulware target can momentarily rise as the time normalisation updates, this ratchet guarantees the acceptance bar never increases, eliminating a class of missed agreements where a good late offer would otherwise be rejected against a transiently higher bar.

5 Opponent Model

The opponent model is an adaptive mixture of two classical frequency estimators over the opponent’s offer stream. The **Smith** component infers, per issue, value scores from observed value frequencies and issue weights from a combination of value *dominance* and *stability* (issues whose value is held constant across consecutive offers are deemed more important). The **distribution-based** component (DFM) uses Laplace-smoothed value distributions, an exponential value filter, and a χ^2 test over disjoint windows to detect genuine concession before increasing the weight of stable issues. The mixture weight on the DFM grows with the number of observations, $w_{\text{dfm}} = 0.35 + 0.30c$ with confidence c rising as more offers arrive, because the windowed DFM becomes reliable only after enough data while Smith is informative from the first rounds. The same model is exposed through the `opponent_ufun` interface and therefore supplies the Concealing numerator a_i in Eq. (1).

6 Evaluation

We evaluate with an ordered head-to-head tournament over all seven official scenarios (CAR, GROCERY, LAPTOP, AMSTERDAM, CAMERA, ISBTACQUISITION, NICEORDIE) with three repeats per pairing, against the official BOANEG and MAPNEG agents and the NegMAS BOULWARE, CONCEDER, ASPIRATION and MICRO baselines. Table 1 reports mean scores; LionelNeg ranks first, ahead of both official baselines and of our own earlier version, though the margin over BOANEG is narrow ($\Delta \approx 0.005$ in total score), close to the noise level of a three-repeat tournament.

Agent	Score	Advantage	Concealing
LionelNeg (ours)	1.478	0.728	0.750
BOANEG (official)	1.473	0.723	0.750
MAPNEG (official)	1.464	0.713	0.751
LionelNeg (prior version)	1.462	0.712	0.750

Table 1: Mean score over seven scenarios, three repeats, head-to-head field. Concealing values are normalised per negotiation, so each agent’s share is below the two-player 0.785 figure seen in smaller fields.

Per scenario, LionelNeg is strongest on the large CAR domain (Advantage ≈ 0.90) and competitive everywhere except the tiny LAPTOP domain (25 rational outcomes), where the Nash-product structure of the frontier limits the utility attainable in the seller role regardless of concession discipline. The gain over our prior version comes almost entirely from Advantage: a sharper Boulware exponent (t^5 instead of t^4), a slightly higher mid-game floor, and a tighter acceptance fraction (0.87). Concealing scores are near-identical across all evaluated agents (0.750–0.751), confirming that Advantage drives all meaningful score differences.

7 Lessons and Suggestions

Our central lesson concerns the Concealing term. Reading the scoring code shows it is the zero-sum ratio of Eq. (1) over Kendall rank correlation, with two levers: lower the opponent’s accuracy of us (the *denominator*) or raise our accuracy of them (the *numerator*). We attempted both extensively — proposal-side deception, iso-utility entropy bidding, anti-stability issue masking and a Bayesian own-model on the denominator side, and a decoupled, early-weighted opponent model on the numerator side — and additionally ran a 50-iteration LLM-driven evolutionary search. *None* moved Concealing by more than measurement noise. The reason is structural:

frequency models recover the global outcome ranking robustly from exactly the high-utility offers that Advantage forces us to make, and a frequency model already encodes concession order implicitly (opponents repeat their best offers), so explicit temporal weighting only adds noise. We conclude that for frequency-based opponents the Concealing term is effectively fixed near its ceiling, and that score competition reduces to Advantage maximisation. A practical suggestion for future editions is to include opponents whose preference models are *not* frequency-based, which would make the Concealing lever genuinely contestable.

A second lesson is domain-dependent: in very small domains the Pareto/Nash geometry, rather than the bidding policy, sets the ceiling on attainable Advantage, so scenario-specific tuning there has sharply diminishing returns.

Conclusions

LionelNeg combines a Boulware concession schedule, an opponent-aware multi-criteria offer selector with a Nash-product cooperation term, a monotone ACNEXT acceptance rule, and a combined Smith + distribution frequency opponent model. In our head-to-head evaluation against BOANEG, MAPNEG, and the standard NegMAS baseline agents, LionelNeg ranks first. Our most useful finding is methodological: the Concealing objective is near a structural ceiling for the current opponent pool, so the reliable path to a higher score is disciplined Advantage maximisation rather than information hiding.