POKERSKILL: A STRONG LLM AGENT IN POKER WITHOUT TRAINING OR SOLVERS

Anonymous authors
Paper under double-blind review

ABSTRACT

Heads-up no-limit Texas hold'em (HUNL) is a long-standing challenge for artificial intelligence. The dominant approach solves for an equilibrium with counterfactual regret minimization, at a cost of millions of core-hours. Large language models (LLMs) hold broad poker knowledge but lose heavily when asked to play directly, and rule-based agents are interpretable and training-free but have a low strategic ceiling. We take a different route: pair a general-purpose LLM with an expert skill library that serves as its action interface. We introduce **PokerSkill**, which computes each game state deterministically, retrieves only the applicable fragments of a layered library written by expert players, and bounds the action space with an attack/defense budget before the model chooses. It trains no parameters and queries no solver at inference time. Against GTOWizard, a state-of-the-art GTO benchmark, PokerSkill reduces the losses of five frontier LLMs by 49–74% with a single library that is never retuned, and in $400{,}000$ hands of live play against Slumbot, the strongest publicly playable solver-based bot, its loss rate is -31 ± 23 mbb/hand, not statistically distinguishable from zero. Our key finding is that rule-based skills alone do not constitute a strong strategy, and LLMs alone cannot play well, but their combination yields an agent that requires neither training nor solver access yet competes with systems built on millions of core-hours of computation. To our knowledge this is the strongest LLM-based poker agent reported to date, and among the strongest open-source poker AI systems. An anonymized code repository is provided in the supplementary material.

1 INTRODUCTION

Poker, particularly Heads-Up No-Limit Texas Hold'em (HUNL), is among the most challenging domains in artificial intelligence (Billings et al., 2002). Its game tree exceeds 10^{164} nodes (Johanson, 2013), requiring agents to handle imperfect information, stochastic outcomes, deception, and multi-street planning in one framework. Solving poker has driven advances in decision-making under uncertainty (Bowling et al., 2015; Moravčík et al., 2017; Brown & Sandholm, 2018; 2019b).

The most successful approaches to HUNL rely on equilibrium solvers built on Counterfactual Regret Minimization (CFR) and its variants (Zinkevich et al., 2007; Brown & Sandholm, 2019a; Li & Huang, 2026c), which traverse the game tree toward a Nash equilibrium. The cost is immense: Libratus required over 15 million core-hours (Brown & Sandholm, 2018), and modern solvers demand dedicated GPU clusters and months of training (Brown et al., 2020). A second tradition is rule-based expert systems (Billings et al., 1998; Davidson et al., 2000; Billings et al., 2002), which encode human poker knowledge as hand-crafted heuristics. They are interpretable and need no training, but rigid rules cannot express the conditional strategies no-limit play requires, and the gap to solver-based systems has been hundreds of mbb/hand (Rubin & Watson, 2011).

LLMs are a third candidate. Frontier models can explain pot odds, minimum defense frequency, and polarized ranges on request. At the table they fail: the GTOWizard benchmark (Provost et al., 2026) reports that Claude Opus 4.6 loses 204 ± 44 mbb/hand under default prompting, and its transcripts show models misreading their own hole cards, mislabeling board texture, and choosing illegal or incoherent bet sizes. Poker knowledge is present in these models, but it is not reliably converted into actions. The three paradigms are summarized in Figure 1.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

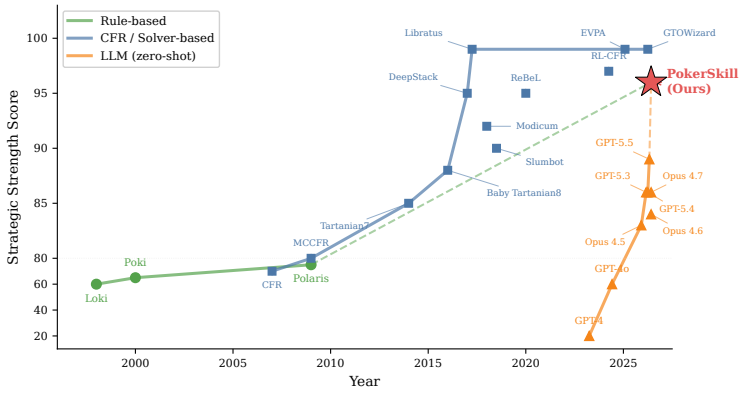


Figure 1: Three paradigms of poker AI. Rule-based agents (green) are interpretable but strategically shallow. CFR/solver-based systems (blue) reach near-equilibrium play through massive computation. LLM agents (orange) hold strategic knowledge but play poorly. PokerSkill (red star) uses expert rule-based skills as the action interface of an LLM and reaches competitive play without training or solver access. The vertical axis is an illustrative ordering, not a measured quantity.

Together they set up the question this paper asks. Equilibrium solving buys strength with computation that only well-resourced groups can afford. Expert rules buy interpretability and zero training with a ceiling that no-limit play exposes. LLMs hold the knowledge but not the reliability. What is missing is a way to make what experts already know act as an interface for a general model: a deterministic scaffold that fixes the state, hands the model only the skills the situation calls for, and narrows the choice to the actions a strong player would consider, so that the model's judgement is spent on the decision rather than on reconstructing the situation.

PokerSkill supplies that scaffold with expert rules, in three stages. First, a *context engine* reads the game state directly and turns board texture, hand class, action line, position, stack-to-pot ratio, and cumulative betting pressure into a handful of compact labels. The model is never asked to reconstruct the state, so the most visible failure of default prompting, misreading one's own cards or the board, is removed from its job. Second, the labels act as a retrieval key into a *layered skill library* written by expert players among the authors: protocol and output format, preflop range tables, stable postflop principles, board- and hand-specific strategy, and river blocker guidance. A decision pulls in only the layers its context calls for, from 60 action-line scenarios, 23 hand classes, and 47 bet-size pressure thresholds; the full library never enters a prompt. Preflop, the library acts directly: the context engine identifies one of eleven preflop scenario families and the action is drawn from the expert range table for that scenario, so the model is consulted preflop only in lines the tables do not cover. Third, an *attack/defense budget* gives each hand class a finite capacity to bet and to defend, spends it as the betting line develops, and collapses the action space to a small set of viable options. The division of labour is therefore explicit: postflop, the rules decide which actions are worth considering and the model decides which one to take. No parameter is updated and no solver is queried.

We evaluate against two opponents. GTOWizard (Provost et al., 2026) is a state-of-the-art GTO benchmark with AIVAT variance reduction (Burch et al., 2018); it beats an earlier release of Slumbot (Jackson, 2013), the strongest publicly available HUNL bot, by 194 ± 41 mbb/hand. With PokerSkill, Claude Fable 5 reaches -44 ± 25 mbb/hand against GTOWizard, GPT-5.5 XHigh -57 ± 21, Claude Opus 4.6 -80 ± 29, DeepSeek-V4-Flash -82 ± 17, and Claude Opus 4.7 -87 ± 33; relative to the default-prompt baselines of the same models (-132 to -204), losses fall by 49–74%. We also play 400,000 hands directly against the 2026 version of Slumbot: DeepSeek-V4-Flash with PokerSkill scores -31 ± 23 mbb/hand, statistically indistinguishable from zero, while the same model without skills loses 203 ± 25. Further ablations show that executing the skill library's chosen actions without an LLM, or removing any one component of PokerSkill, costs performance. The expert library alone is therefore of limited strength and an LLM alone plays weakly, but combining the two yields more than the sum of its parts.

Our contributions are:

- **A skill-grounded poker agent.** PokerSkill is training-free and solver-free. A deterministic context engine extracts the features of each decision, retrieves the matching entries of an expert-written skill library, and bounds the action space; an unmodified LLM then chooses within those bounds. To our knowledge it is the strongest LLM-based poker agent reported to date, and among the strongest open-source poker AI systems.
- **An evaluation that separates the ingredients.** Five frontier LLMs from three providers improve by 49–74% under rules that are never retuned; the skill library executed with no LLM at all reaches -132 ± 19 mbb/hand, on par with the best default-prompted model; and a leave-one-out ablation shows that removing context-specific strategy alone returns the agent to baseline. The gain holds under two opponents for DeepSeek-V4-Flash, the model run against both: the GTOWizard benchmark and the live 2026 Slumbot server, both reported in mbb/hand.
- **A transferable design pattern.** When a domain is too complex to solve directly and too costly to learn from scratch, expert knowledge can be written down as a precise, deterministic scaffold that constrains and supports a general-purpose model. Here the scaffold is a state abstraction, a retrieval key, and a bounded action set, and a general-purpose language model inside it reaches a level of play in a specialised, uncertain, adversarial domain that otherwise requires solver-based computation, at no task-specific training cost and no external search at inference time. We demonstrate and quantify the pattern in HUNL.

2 Related Work

Poker solvers and game-theoretic agents. CFR and its variants (Zinkevich et al., 2007; Lanctot et al., 2009; Brown & Sandholm, 2019a; Farina et al., 2021; Li & Huang, 2026b; Li et al., 2026b; Li & Huang, 2026c) are the foundation for solving large imperfect-information games. Superhuman HUNL systems combine CFR with abstraction (Ganzfried & Sandholm, 2014; Li et al., 2024; Li & Huang, 2025; 2026d), depth-limited solving (Brown et al., 2018; 2020), and subgame resolving (Brown & Sandholm, 2017): Libratus (Brown & Sandholm, 2018) required over 15 million core-hours, and DeepStack (Moravčík et al., 2017) generated its training data with computation measured in core-years. Slumbot (Jackson, 2013) remains the strongest publicly playable HUNL bot, and GTOWizard (Provost et al., 2026) is the current state-of-the-art benchmark, beating an earlier release of Slumbot by 194 ± 41 mbb/hand.

LLM poker agents. LLMs can discuss poker strategy but struggle to play reliably. Gupta (Gupta, 2023) found ChatGPT and GPT-4 deviate from GTO strategies. Suspicion-Agent (Guo et al., 2024) uses theory-of-mind prompting but only in small Leduc Hold'em. PokerBench (Zhuang et al., 2025) evaluates LLM poker reasoning; even fine-tuned models underperform solvers. ToolPoker (Lin et al., 2026) queries external solvers for GTO actions in Limit Hold'em. These approaches restrict the game variant, fine-tune the model, or leave the strategy to a solver (Li & Huang, 2026a); PokerSkill targets full HUNL with an unmodified LLM inside a deterministic expert scaffold and needs none of these.

LLM agents in strategic settings. Cicero (Meta Fundamental AI Research Diplomacy Team (FAIR) et al., 2022) achieved human-level Diplomacy by combining language models with strategic planning. Voyager (Wang et al., 2023) uses LLMs with a skill library for open-ended embodied tasks, an architectural parallel to PokerSkill's layered retrieval. In social deduction games, LLM agents exhibit emergent strategic behaviors but remain suboptimal (Bailis et al., 2024), and raw LLM capabilities generally require structured support for effective situated action (Xi et al., 2025).

Rule-based poker agents. The earliest computer poker programs relied on hand-crafted heuristics. Loki (Billings et al., 1998) and its successor Poki (Davidson et al., 2000) combined rule-based betting logic with opponent modeling; later work explored game-theoretic approximations within rule-based frameworks (Billings et al., 2003), and Polaris (Bowling et al., 2009) combined rule-based components with equilibrium computation. Despite their interpretability and zero training cost, rule-based agents hit a strategic ceiling well below equilibrium play (Rubin & Watson, 2011). PokerSkill revisits this approach. Its rule-based skills alone play at -132 mbb/hand against GTOWizard (Section 4); as the action interface of a modern LLM, which supplies the contextual judgment that rigid rules lack, they

yield competitive HUNL play without training or solver access. Appendix F relates PokerSkill to retrieval-augmented and structured-prompting methods.

3 POKERSKILL

PokerSkill wraps an unmodified LLM in a deterministic layer of expert poker rules (Figure 2). Its purpose is concrete: default-prompted LLMs fail at the table in three ways, misreading the game state, applying a strategic principle that does not fit the situation, and choosing invalid or incoherent bet sizes. In the GTOWizard transcripts (Provost et al., 2026), Claude Opus 4.6 holding 4♡7♠ on the turn 5♡4♣3♡4♠ states "I have QKo on a 5d4c3h4s board, complete air with no pair" while actually holding trips. PokerSkill is built to remove each failure by computing the state deterministically, retrieving only the rules that apply, and bounding the action space before the model decides.

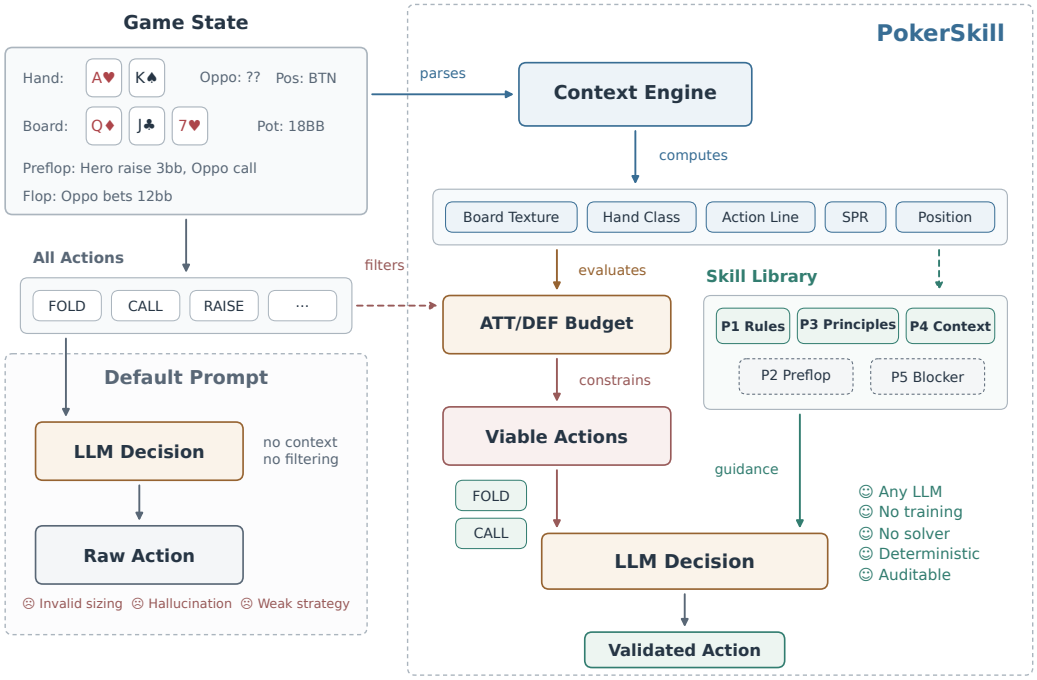


Figure 2: PokerSkill versus the default prompting baseline. **Left**: a bare LLM receives the game state and all legal actions, producing raw outputs prone to invalid sizing, hallucination, and weak strategy. **Right**: PokerSkill interposes a deterministic Context Engine that extracts structured features (board texture, hand class, action line, SPR, position), an ATT/DEF Budget mechanism that filters the action space into viable options, and a layered Skill Library that provides situation-specific guidance. The LLM decides within these bounds, and a validator ensures the final output is a legal action.

Each postflop decision passes through three stages: (i) **Context Engine**: deterministic rules analyze cards, board texture, action history, position, stack-to-pot ratio (SPR), and cumulative betting pressure, producing compact labels. (ii) **Skill Retrieval**: the labels select specific prompt fragments from a layered skill library; the full library never enters a prompt. (iii) **Bounded Decision**: the LLM chooses among context-compatible actions within validated legal bounds. Preflop decisions follow the range tables of Section 3.2 directly.

The entire skill library, including prompt fragments, budget tables, context labels, and action constraints, is written by experienced poker players among the authors who have studied GTO theory and played at high stakes. It covers 60 action-line scenarios, 23 hand classes, and 47 bet-size pressure thresholds, and the same tables are used unchanged for every LLM here. "Solver-free" means that no equilibrium solver is invoked at inference time: the library is informed by GTO theory and solver study, as human professionals are, but it performs no game-tree traversal or iterative computation.

3.1 THE CONTEXT ENGINE

The context engine computes features in four groups. **Board features** are dry/wet texture, paired, trips and monotone boards, flush-completing and straight-completing boards, and one-card straight or flush threats. **Hand features** are made-hand class, draw class, kicker class, and high-card showdown labels. **Line features** are pot type, initiative, position, street, prior bet count, weighted bet pressure, and scenario ID. **Stack features** are SPR and all-in thresholds.

Board analysis and special boards. Board classification covers texture, suit structure, and rank structure; Table 4 in Appendix C.4 lists the labels and their effect on the prompt. Paired, trips, double-paired, quads and one-card-straight or one-card-flush boards are classified separately, because on those boards ordinary hand ordering misleads and a nominally strong hand can be dominated or completed by a single private card. Each special board type receives a complete strategy override with its own pre-specified attack/defense budgets rather than an adjustment to the normal tables. Appendix C.4 gives the full label set, the failure each one prevents, and the override budgets.

Hand analysis. The hand analyzer returns a made-hand label (15 classes, from `nuts` down to no-showdown trash) and a draw label (8 classes by draw quality and overcard rank). Together with the line and stack features these labels select the P4 fragments and the budget row for the current decision; Appendix D lists all 23 with their budgets.

3.2 CONTEXT-CONDITIONED SKILL RETRIEVAL

The skill library is organized into five layers, each activated by different context conditions:

- **P1** (always active): Game rules, legal actions, output format constraints.
- **P2** (range table): Range guidance for the current preflop scenario.
- **P3** (general principles): Stable principles including value/bluff separation, positional strategy, pot control, and sizing discipline.
- **P4** (context-specific): Strategy keyed by board texture, hand class, action line, and role.
- **P5** (blocker): Blocker-aware bluffing and bluff-catching guidance for river decisions.

A preflop decision in one of the eleven covered scenario families is played from the P2 range table without a model call; in uncovered lines the model receives P1 + the range text for the nearest scenario. A postflop decision retrieves P1 + P3 + P4 entries matching the board texture and hand class. A river decision additionally retrieves P5. Selective retrieval avoids vague general advice that does not resolve the decision and an overload of irrelevant heuristics to arbitrate among.

Preflop range tables (P2). The preflop game tree decomposes into a finite set of well-studied scenarios (e.g., button open, big blind facing a 3-bet). We encode 12 scenario families at 200BB depth, from initial open/limp decisions through 5-bet and all-in confrontations, each with frequency-annotated ranges over the 169 starting-hand classes authored by the expert authors. At decision time the context engine identifies the scenario; for the eleven covered families the agent samples its action from the table's mixed strategy and does not call the model, and in the residual family the model is prompted with the range text. Preflop is therefore where the library plays alone and postflop is where the model's judgement enters.

Stable postflop principles (P3). P3 encodes principles that competent players apply but that LLMs frequently violate under default prompting: (1) *value/bluff separation*, polarizing bets into strong hands and bluffs while medium-strength hands prefer to check; (2) *positional strategy*, adapting aggression and defense thresholds to acting first or last; (3) *pot control*, not inflating the pot with vulnerable holdings; and (4) *sizing discipline*, matching bet sizes to hand polarity and board texture. P3 is included in every postflop prompt as the strategic floor that P4 and P5 build upon.

Action-line scenario recognition (P4 action lines). The action-line analyzer turns betting history into a scenario ID and a cumulative pressure estimate. Scenario IDs distinguish whether the hero is the aggressor or defender, whether the opponent has checked or bet, and whether the line is a c-bet,

probe, donk, check-raise, or others. We define 60 scenarios across flop (9), turn (21), and river (30). For example, scenario `R-AA1` ("we bet flop and turn, opponent called both, checks river") marks the opponent's range as likely *capped*, which is when a polarized river bet is worth considering; scenario `T-D2` ("opponent bet both flop and turn") signals sustained aggression, and the defense thresholds tighten accordingly. Weighted pressure is computed from bet size relative to the pot: small bets consume little budget, pot-sized or overbet actions much more, so

$$\text{larger bet fraction} \Rightarrow \text{larger pressure weight} \Rightarrow \text{less remaining ATT/DEF}.$$

Compressing the multi-street history into a semantic label gives the LLM the *conclusion* in one line.

Board texture classification (P4 board texture). A dry board usually favors the preflop aggressor's range and supports high-frequency small bets; a wet board such as $J\heartsuit T\heartsuit 8\clubsuit$ distributes equity more evenly and demands larger, less frequent bets. Models mislabel texture (e.g., calling a two-tone connected board "dry"), which cascades into wrong strategy; the deterministic label fixes the premise.

Hand strength classification (P4 hand strength). The same hand (e.g., top pair with a king kicker) can be a strong value hand on a dry board in a single-raised pot or a marginal bluff-catcher on a wet board in a 3-bet pot. We define 15 made-hand labels and 8 draw labels, and a hand carries one of each. The *made-hand system* (15 classes) assigns ATT/DEF budgets based on showdown value: strong hands receive high budgets for multi-street aggression, marginal hands limited budgets that confine them to checking or single-street defense. The *drawing-hand system* (8 classes) gives draws ATT budgets for semi-bluffing but governs their defense by *pot-odds thresholds* specifying the maximum callable bet size per street, because draw equity depends on remaining cards rather than accumulated pressure. A hand with both made-hand value and a draw (e.g., a pair plus a flush draw) follows a *combo rule* that adds the draw's contribution to the made-hand baseline.

River blocker guidance (P5 Blocker). On the river, showdown value is fixed and *blocker effects* become a first-order factor: holding a card that reduces the opponent's strong combinations makes bluffing more profitable, while *not* blocking the opponent's bluffs makes calling more profitable. P5 provides river-specific guidance on (1) *when to bluff*: hands with no showdown value that block the opponent's value range and keep a consistent betting narrative; and (2) *when to bluff-catch*: hands with showdown value that do not block the opponent's bluffs, with minimum defense frequency as a range-level reference. It is activated only on the river, where draw equity no longer dominates.

3.3 Attack/Defense Budget System

The budget system translates context into action constraints, following the principle that *hands have finite strategic capacity*: a top-pair hand can typically bet two streets and defend all three, while a middle pair can defend one or two streets before folding becomes correct. Because this is the part of the interface that removes actions, we state it precisely.

Definition 1 (Pressure and remaining budgets)**.** *Let h be the hero's observable state (private cards, board, action history, stacks) and $c = \phi(h)$ its context labels. Index the postflop bets and raises of* either *player so far in the hand by $i = 1, \ldots, m$; let r_i be the incremental amount of bet i and P_i the pot immediately before it. Its pressure weight is $w_i = \omega(r_i/P_i)$, where ω is a fixed non-decreasing step function with 47 steps, from 0.04 below 5% of the pot to 4.0 at 1500% and above; calls and preflop actions carry no weight. With $W(h) = \sum_{i \le m} w_i$ and w_{cur} the weight of a bet the hero currently faces (zero if none),*

$$B^{\text{att}}_{\text{rem}}(h) = B^{\text{att}}(c) - W(h), \qquad B^{\text{def}}_{\text{rem}}(h) = B^{\text{def}}(c) - W(h) + w_{\text{cur}}, \tag{1}$$

where the base budgets $B^{\text{att}}(c), B^{\text{def}}(c)$ are expert-authored tables indexed by the made-hand and draw labels, pot type, board labels, kicker, position, and street, with special-board overrides replacing both values (Appendix D). Both quantities are recomputed from the action history at every decision; nothing is carried between decisions.

The budget is a heuristic encoding of three ideas from equilibrium play, without computing an equilibrium: pressure is spread across streets, continuation is rationed rather than unlimited, and low-SPR spots reduce to commit-or-fold. Base budgets range from a value that no betting line can exhaust (the nuts) down to zero for hands with no betting purpose, and context modifiers lower them

on wet or flush-possible boards, for weaker kickers, in larger pot types, and out of position. Drawing hands defend by pot-odds thresholds instead of the defense budget. Because W grows with every bet, a hand that bets the flop has less budget left for the turn, which carries earlier pressure into later decisions, a bookkeeping that default prompting has no equivalent of.

How the budgets were calibrated. The ATT/DEF values are set by human poker experts among the authors, drawing on years of professional play and coaching experience. For each hand category and board texture, the experts specify how aggressively (ATT) or defensively (DEF) a hand should be played across multiple streets. Top pair on a dry board, for example, receives ATT ≥ 1.0 because expert consensus is that such hands should bet at least one street, while marginal hands receive low ATT budgets because they should mostly check. The numeric values are calibrated so that the weighted-pressure cost of standard betting lines (e.g., 33% flop + 66% turn + 100% river) aligns with expert intuition about how many streets each hand class should bet. The same tables improve five different LLMs consistently, so they encode strategic content that is not specific to a model.

From budget to viable actions. The remaining budgets become an action set as follows.

Definition 2 (Viable options)**.** *Let $\mathcal{A}$ be the abstract action vocabulary:* FOLD, CHECK, CALL*, and* BET *or* RAISE *with a size class. The viable-option map $V(h) \subseteq \mathcal{A}$ is a deterministic function of $B^{\text{att}}_{\text{rem}}$, $B^{\text{def}}_{\text{rem}}$, w_{cur}, street, position, SPR, and c, evaluated as an ordered cascade: (i) a drawing hand facing a bet on the flop or turn is judged by a pot-odds threshold for its draw class; (ii) below an SPR threshold (*2.0 *out of position,* 1.2 *in position) a hand with showdown value and positive remaining attack budget is committed, so* FOLD *is removed and, once $B^{\text{att}}_{\text{rem}} \geq \text{SPR}$, only an all-in or a geometric bet remains; (iii) facing a bet,* RAISE *requires $B^{\text{att}}_{\text{rem}}$ above a street- and position-specific threshold (from 0.75 on the river in position to 3.0 on the flop out of position at high SPR),* CALL *alone is offered when the defense budget left after calling, $B^{\text{def}}_{\text{rem}} - w_{\text{cur}}$, exceeds 0.5,* FOLD *alone when it is negative, and both when it lies in $[0, 0.5]$; (iv) not facing a bet, a value bet requires $B^{\text{att}}_{\text{rem}} \geq 1$, with sizes set by street, texture, and role, and otherwise a check or a small bet is offered, with paired-board overrides removing thin value bets. A value within 0.12 of a threshold unlocks the option on both sides. Every branch of the cascade returns a non-empty set. $V(h)$ enters the prompt as text with sizing ranges; it is guidance, and legality is enforced separately (Proposition 1).*

Action grounding. The final prompt contains the game state, the selected skill entries (P1–P5 fragments relevant to this decision), the computed context verdicts, the viable actions with sizing ranges, and a structured output schema. The LLM returns structured JSON via forced tool-use, and a validator ensures the action is legal and the size falls within bounds. If the response cannot be parsed, the agent checks if legal, otherwise calls, otherwise folds; an out-of-range size is clamped to the legal range. Apart from the LLM call and the sampling of mixed preflop actions from the range table, the system is deterministic: given the same game state and the same model response, the prompt, budgets, viable options, and validated action are reproduced exactly.

Proposition 1 (Legality and totality)**.** *Let $L(h) \neq \emptyset$ be the legal action set supplied by the environment. The validator maps every model response to an action in $L(h)$: a parsed action outside $L(h)$ is replaced through a fixed substitution table that prefers the passive legal action, a size outside the legal range is clamped to it, and an unparseable response becomes* CHECK *if legal, else* CALL*, else* FOLD*. Hence the executed action is legal at every decision, for every possible model output. Membership in $V(h)$ is not enforced.*

Proposition 2 (Street capacity of a hand class)**.** *Fix the context labels c, so that $B^{\text{att}}(c) = B$ is fixed, and suppose only the hero bets, at pot fractions $f_1, f_2, \ldots$ on successive streets. Since $\omega \geq 0.04$, $B^{\text{att}}_{\text{rem}}$ is strictly decreasing in the number of bets, and rule (iv) offers a value bet on the k-th street only when $B - \sum_{j<k} \omega(f_j) \geq 0.88$ (the preflop aggressor's flop c-bet and the closing river bet in position have their own, lower thresholds). The statement is bookkeeping: it says nothing about whether the offered actions are correct, and it fails when the labels change between streets, since the base budget is then re-read from the tables.*

These statements are the whole of what the interface guarantees: the same labels for the same state, a legal executed action, and the budget arithmetic above. Whether the offered actions are the right ones is the question Section 4 answers.

4 EXPERIMENTS

We evaluate five frontier LLMs and two baselines, each model under a default prompt and the skill library executed with no LLM, against two opponents: GTOWizard (Provost et al., 2026), a state-of-the-art GTO benchmark that reports AIVAT variance-reduced (Burch et al., 2018) results, and the 2026 version of Slumbot (Jackson, 2013), the strongest publicly playable HUNL bot. Every configuration plays at least 5,000 hands; the two Slumbot runs play 400,000 and 164,700 hands. Appendix E gives the model settings and the full protocol of both evaluations, and Appendix A gives the unit conversions.

Main results. Table 1 and Figure 3 report the outcome. PokerSkill lowers the loss of every model it is applied to, by between 49% and 74%, and the gain is more than two standard errors from zero for four of the five ($z = 2.3$–3.2). Claude Opus 4.7 is the exception at $z = 1.9$, though its point estimate improves by the same order as the others. No default-prompted model reaches -130 mbb/hand, and all five PokerSkill agents beat -90.

What the rules and the model each contribute. Neither half is strong alone. The skill library executed with a fixed first-viable-action selector and no model at all reaches -132 ± 19 mbb/hand, on par with the best default-prompted LLM and better than the other four, and default prompting with no rules lands between -132 and -204. Together they reach -44 to -87. The gap between the rule-based agent and each PokerSkill agent measures what replacing a fixed selector by the model is worth: 50–88 mbb/hand for Claude Fable 5, GPT-5.5, and DeepSeek-V4-Flash ($z = 2.0$–2.8), with both Opus models inside noise. The budget system constructs the actions a strong player would consider; it does not rank them. Choosing among them is where the language model earns its place, and it is the part these rules do not attempt: the hand-written strategy alone does not rise above a default-prompted model, and no model reaches the PokerSkill range without it. The combination trains no parameters, against the millions of core-hours that equilibrium solving consumes offline.

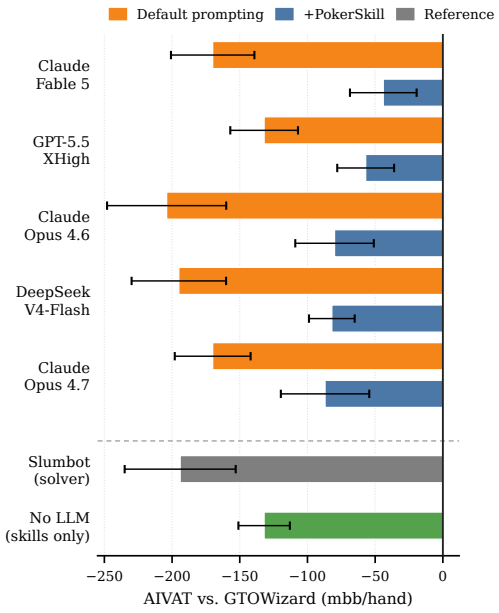


Figure 3: Default prompting versus PokerSkill against GTOWizard (AIVAT mbb/hand, higher is better, error bars = 1 SE; the Slumbot reference is quoted as reported in (Provost et al., 2026)). Rows below the divider are references; values in Table 1.

Cross-model consistency. The same tables, unchanged, improve five models from three providers whose default-prompted losses span 72 mbb/hand. Claude Opus 4.6 starts worst at -204 and ends at -80, and the open-weight DeepSeek-V4-Flash lands at -82 with the tightest interval. After PokerSkill the five span -44 to -87, and no pairwise difference among them is resolvable at these sample sizes. Applying the interface also reorders them: GPT-5.5 XHigh is the strongest default-prompted agent and drops to second, and Claude Opus 4.6 moves from last to third, so the point-estimate ranking before the interface is applied is not preserved inside it, although the post-interface differences are within noise.

Direct play against Slumbot. GTOWizard supplies AIVAT and a fixed reference strategy; Slumbot supplies neither, so the second evaluation trades statistical efficiency for realism. Over 400,000 hands, DeepSeek-V4-Flash with PokerSkill scores -31 ± 23 mbb/hand against the 2026 Slumbot, statistically indistinguishable from zero ($z = 1.3$); the sample does not resolve whether the agent is ahead of or behind the bot. The same model under default prompting loses 203 ± 25 mbb/hand, so the skill library is worth 172 ± 34 mbb/hand here ($z = 5.1$).

Table 1: AIVAT performance against GTOWizard (mbb/hand, ± = 1 SE, higher is better). †Quoted from (Provost et al., 2026).

Agent	Default	+PokerSkill
Claude Fable 5	-170 ± 31	-44 ± 25
GPT-5.5 XHigh	-132 ± 25	-57 ± 21
Claude Opus 4.6	$-204 \pm 44^{\dagger}$	-80 ± 29
DeepSeek-V4-Flash	-195 ± 35	-82 ± 17
Claude Opus 4.7	-170 ± 28	-87 ± 33
Rule-based (no LLM)		-132 ± 19
Slumbot	-194 ± 41	

Table 2: Leave-one-out ablation with DeepSeek-V4-Flash against GTOWizard (mbb/hand, ± = 1 SE, higher is better).

Configuration	mbb/hand
Full PokerSkill	-82 ± 17
w/o P1 system rules	-135 ± 33
w/o P2 preflop ranges	-109 ± 16
w/o P3 postflop principles	-110 ± 24
w/o P4 context-specific	-201 ± 32
w/o P5 river blockers	-129 ± 22
w/o viable-option filter	-146 ± 26
Default prompt	-195 ± 35

Cross-opponent agreement. The two opponents measure the skill effect in different ways: AIVAT-adjusted mbb/hand against a fixed reference strategy, and raw chips against a live server. For DeepSeek-V4-Flash, the only model run against both, the skill effect is 113 ± 39 mbb/hand under AIVAT and 172 ± 34 mbb/hand in live play, and the two estimates differ by 59 ± 51 mbb/hand ($z = 1.1$). An opponent that does not use GTOWizard's reference strategy therefore yields a gain of the same order. The two opponents cannot be compared directly, because GTOWizard's -194 mbb/hand line for Slumbot was measured against an earlier release of that bot.

Ablation of the skill layers. Table 2 removes one layer at a time. Context-specific strategy (P4) is the layer the system depends on: without it the agent falls to -201 ± 32 mbb/hand, a loss of 120 ± 36 relative to the full system and a point estimate close to the default prompt (-195 ± 35; difference 6 ± 47), so the data do not resolve what the other layers achieve without it. The second largest effect is the viable-option filter, worth 64 ± 31 mbb/hand ($z = 2.0$); it is not a skill layer but the constraint that keeps the action space strategically coherent. The four remaining layers each cost between 27 and 53 mbb/hand with overlapping intervals ($z = 1.0$–1.7), and this experiment does not rank them; removing P2 hands preflop play to the model with the game state alone and costs 27 ± 23 mbb/hand. Summed, the six leave-one-out deltas reach 340 mbb/hand against a total gain of 113 over the default prompt; the excess of 227 ± 112 mbb/hand ($z = 2.0$, Appendix E.4) indicates that the layers are not additive contributions. The picture consistent with these numbers is a context-specific core wrapped in general principles and action constraints.

5 Discussion and Conclusion

This paper is the first to show that a general-purpose language model, given an expert skill library as its action interface, becomes a strong heads-up no-limit hold'em agent with no training and no solver. Expert knowledge written once lifts five frontier models to a level of play that approaches solver-based, near-equilibrium opponents, and it does so under a new paradigm: the model is not trained to play poker, it is interfaced to what experts already know. Strength in one of the hardest games of imperfect information no longer has to be bought with millions of core-hours.

The broader lesson is about the interface rather than about poker. The scaffold has three parts, and each answers a distinct failure of the bare model. The context engine fixes the state, so the model never has to reconstruct its own cards or the board; retrieval hands it only the skills the situation calls for; and the budget system leaves it a choice among a few viable options, so its judgement is spent where a strong player's is. None of these parts is a poker algorithm. The library contents are the only poker-specific ingredient, written once by expert players and reused across five models and two opponents without retuning. The ablation makes the same point from the other side: no single layer carries the result, the rules without the model stop at the level of a default-prompted LLM, and the model without the rules plays badly, so the gain lives in the combination and not in either half. Wherever expert knowledge can be written down this way, we expect the same division of labour to carry over, and to improve as the models do.

AI USE STATEMENT

We used Claude Code (Anthropic) to assist with experiment scripting, data analysis, figure generation, and LaTeX formatting. LLM-generated code was reviewed and tested by the authors, and every number in the paper is recomputed from the raw hand logs recorded during the runs. Generative AI tools were not used to generate research ideas, to write the skill library, or to conduct the literature survey; all scientific claims, experimental designs, and interpretations are the authors' own. The LLMs evaluated in this paper (GPT-5.5, Claude Opus 4.6/4.7, Claude Fable 5, DeepSeek-V4-Flash) are the objects of study.

ETHICS STATEMENT

This work studies LLM-based poker agents in a controlled research setting against the GTOWizard benchmark and the public Slumbot server, both of which are designed for bot-versus-bot evaluation. We do not deploy these agents on real-money platforms or against human players. All external services (GTOWizard, Slumbot, OpenAI API, Anthropic API, DeepSeek API) were used under their respective terms of service for research purposes. Poker AI could in principle be misused for online gambling; our agents operate at 15–20 seconds per decision with frontier models, play only two-player 200BB games, and major poker platforms already employ bot detection. We encourage use of this research for advancing AI reasoning.

REPRODUCIBILITY STATEMENT

Apart from the LLM call and the sampling of preflop actions from the range table, the system is deterministic: the context engine, skill library, budget tables, and validator are fully specified in Section 3 and Appendix D, and complete prompt traces are given in Appendix G. The evaluation protocol (hand counts, AIVAT, model settings, dates) is in Section 4 and Appendix E. An anonymized code package is provided as supplementary material: it exposes the same entry points we used to run every experiment in this paper, with the core decision logic shipped as compiled extensions, and it can be pointed at GTOWizard or Slumbot to regenerate results of the same form. The full source will be open-sourced upon acceptance.

REFERENCES

Suma Bailis, Jane Friedhoff, and Feiyang Chen. Werewolf arena: A case study in llm evaluation via social deduction. *arXiv preprint arXiv:2407.13943*, 2024.

Darse Billings, Denis Papp, Jonathan Schaeffer, and Duane Szafron. Opponent modeling in poker. In *AAAI/IAAI*, pp. 493–499. AAAI Press, 1998.

Darse Billings, Aaron Davidson, Jonathan Schaeffer, and Duane Szafron. The challenge of poker. *Artificial Intelligence*, 134(1-2):201–240, 2002.

Darse Billings, Neil Burch, Aaron Davidson, Robert Holte, Jonathan Schaeffer, Terence Schauenberg, and Duane Szafron. Approximating game-theoretic optimal strategies for full-scale poker. In *International Joint Conference on Artificial Intelligence*, pp. 661–668, 2003.

Michael Bowling, Neil Burch, Michael Johanson, and Oskari Tammelin. Heads-up limit hold'em poker is solved. *Science*, 347(6218):145–149, 2015.

Michael H. Bowling, Nicholas Abou Risk, Nolan Bard, Darse Billings, Neil Burch, Joshua Davidson, John Alexander Hawkin, Robert Holte, Michael Johanson, Morgan Kan, Bryce Paradis, Jonathan Schaeffer, David Schnizlein, Duane Szafron, Kevin Waugh, and Martin Zinkevich. A demonstration of the Polaris poker system. In *AAMAS*, pp. 1391–1392, 2009.

Noam Brown and Tuomas Sandholm. Safe and nested subgame solving for imperfect-information games. *Advances in neural information processing systems*, 30, 2017.

Noam Brown and Tuomas Sandholm. Superhuman ai for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374):418–424, 2018.

Noam Brown and Tuomas Sandholm. Solving imperfect-information games via discounted regret minimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 1829–1836, 2019a.

Noam Brown and Tuomas Sandholm. Superhuman ai for multiplayer poker. *Science*, 365(6456): 885–890, 2019b.

Noam Brown, Tuomas Sandholm, and Brandon Amos. Depth-limited solving for imperfect-information games. *Advances in neural information processing systems*, 31, 2018.

Noam Brown, Anton Bakhtin, Adam Lerer, and Qucheng Gong. Combining deep reinforcement learning and search for imperfect-information games. *Advances in neural information processing systems*, 33:17057–17069, 2020.

Neil Burch, Michael Johanson, and Michael Bowling. AIVAT: A new variance reduction technique for agent evaluation in imperfect information games. In *AAAI Conference on Artificial Intelligence*, 2018.

Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W. Cohen. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856.

Aaron Davidson, Darse Billings, Jonathan Schaeffer, and Duane Szafron. Improved opponent modeling in poker. In *International Conference on Artificial Intelligence (ICAI)*, pp. 1467–1473, 2000.

Gabriele Farina, Christian Kroer, and Tuomas Sandholm. Faster game solving via predictive blackwell approachability: Connecting regret matching and mirror descent. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 5363–5371, 2021.

Sam Ganzfried and Tuomas Sandholm. Potential-aware imperfect-recall abstraction with earth mover's distance in imperfect-information games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 28, 2014.

Jiaxian Guo, Bo Yang, Paul Yoo, Bill Yuchen Lin, Yusuke Iwasawa, and Yutaka Matsuo. Suspicion agent: Playing imperfect information games with theory of mind aware GPT-4. In *First Conference on Language Modeling*, 2024.

Akshat Gupta. Are chatgpt and GPT-4 good poker players? - A pre-flop analysis. *CoRR*, abs/2308.12466, 2023.

Wenyue Hua, Ollie Liu, Lingyao Li, Alfonso Amayuelas, Julie Chen, Lucas Jiang, Mingyu Jin, Lizhou Fan, Fei Sun, William Wang, Xintong Wang, and Yongfeng Zhang. Game-theoretic LLM: agent workflow for negotiation games. *CoRR*, abs/2411.05990, 2024.

Jen-tse Huang, Eric John Li, Man Ho Lam, Tian Liang, Wenxuan Wang, Youliang Yuan, Wenxiang Jiao, Xing Wang, Zhaopeng Tu, and Michael R Lyu. How far are we on the decision-making of llms? evaluating llms' gaming ability in multi-agent environments. *arXiv preprint arXiv:2403.11807*, 2024.

Eric Jackson. Slumbot nl: Solving large games with counterfactual regret minimization using sampling and distributed processing. In *AAAI Workshop on Computer Poker and Incomplete Information*, 2013.

Michael Johanson. Measuring the size of large no-limit poker games. *arXiv preprint arXiv:1302.7008*, 2013.

Marc Lanctot, Kevin Waugh, Martin Zinkevich, and Michael Bowling. Monte carlo sampling for regret minimization in extensive games. In *Advances in Neural Information Processing Systems*, volume 22, 2009.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in neural information processing systems*, 33: 9459–9474, 2020.

Boning Li and Longbo Huang. Efficient online pruning and abstraction for imperfect information extensive-form games. In *The Thirteenth International Conference on Learning Representations*, 2025.

Boning Li and Longbo Huang. Abstraction agent: Llm-guided abstraction for imperfect-information games. In *Findings of the Empirical Methods in Natural Language Processing*, 2026a.

Boning Li and Longbo Huang. Gpu-cfr: 80x faster counterfactual regret minimization by compiling the game to static dataflow and cuda graph replay. *arXiv preprint arXiv:2609.11923*, 2026b.

Boning Li and Longbo Huang. Real-time parallel counterfactual regret minimization. *arXiv preprint arXiv:2605.19928*, 2026c.

Boning Li and Longbo Huang. Effective, efficient, and general information abstraction for imperfect-information extensive-form games. *arXiv preprint arXiv:2605.10900*, 2026d.

Boning Li, Zhixuan Fang, and Longbo Huang. Rl-cfr: improving action abstraction for imperfect information extensive-form games with reinforcement learning. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 27752–27770, 2024.

Boning Li, Yu Chen, and Longbo Huang. Av-aivat: 74x cheaper agent evaluation with certified anytime-valid stopping in imperfect-information games. *arXiv preprint arXiv:2608.06362*, 2026a.

Boning Li, Yu Chen, and Longbo Huang. Correlated chance sampling for monte carlo counterfactual regret minimization. *arXiv preprint arXiv:2607.27035*, 2026b.

Minhua Lin, Enyan Dai, Hui Liu, Xianfeng Tang, Yuliang Yan, Zhenwei Dai, Jingying Zeng, Zhiwei Zhang, Fali Wang, Hongcheng Gao, Chen Luo, Xiang Zhang, Qi He, and Suhang Wang. How far are LLMs from professional poker players? revisiting game-theoretic reasoning with agentic tool use. In *The Fourteenth International Conference on Learning Representations*, 2026.

Meta Fundamental AI Research Diplomacy Team (FAIR), Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, et al. Human-level play in the game of diplomacy by combining language models with strategic reasoning. *Science*, 378(6624):1067–1074, 2022.

Matěj Moravčík, Martin Schmid, Neil Burch, Viliam Lisý, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh, Michael Johanson, and Michael Bowling. Deepstack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, 356(6337):508–513, 2017.

Marc-Antoine Provost, Nejc Ilenic, Christopher Solinas, and Philippe Beardsell. Gto wizard benchmark. *arXiv preprint arXiv:2603.23660*, 2026.

Jonathan Rubin and Ian D. Watson. Computer poker: A review. *Artificial Intelligence*, 175(5-6): 958–987, 2011.

Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Eric Hambro, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. *Advances in neural information processing systems*, 36:68539–68551, 2023.

Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*, 2023.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.

Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, et al. The rise and potential of large language model based agents: A survey. *Science China Information Sciences*, 68(2):121101, 2025.

Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems*, 36:11809–11822, 2023a.

Shunyu Yao, Jeffrey Zhao, Dian Yu, Nan Du, Izhak Shafran, Karthik Narasimhan, and Yuan Cao. React: Synergizing reasoning and acting in language models. In *International Conference on Learning Representations (ICLR)*, 2023b.

Richard Zhuang, Akshat Gupta, Richard Yang, Aniket Rahane, Zhengyu Li, and Gopala Anumanchipalli. Pokerbench: Training large language models to become professional poker players. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 26175–26182, 2025.

Martin Zinkevich, Michael Johanson, Michael Bowling, and Carmelo Piccione. Regret minimization in games with incomplete information. In *Advances in Neural Information Processing Systems*, volume 20, 2007.

A HUNL RULES AND TERMINOLOGY

Heads-up No-Limit Texas Hold'em (HUNL) is a two-player poker game. Each player receives two private hole cards. Five public board cards are revealed over four betting streets: preflop, flop, turn, and river. The small blind posts half a big blind and acts first preflop; the big blind posts one big blind. After the flop, the button is in position (IP) and acts last, while the big blind is out of position (OOP) and acts first. Positions alternate every hand: the player who was the button becomes the big blind in the next hand, and vice versa.

Legal actions are *fold*, *check*, *call*, *bet*, *raise*, and *all-in*. Checking is legal only when facing no bet. Calling matches the current bet. Betting or raising requires a legal amount within the environment's minimum and maximum bounds. No-limit means the amount can vary up to the effective stack.

We report all poker outcomes in milli-big-blinds per hand (mbb/hand), the unit in which GTOWizard reports its AIVAT-adjusted results; 1 bb/100 equals 10 mbb/hand. AIVAT is a variance-reduction estimator for poker evaluation (Burch et al., 2018). Slumbot does not provide it, so Slumbot outcomes are raw chip results converted to mbb/hand, with more per-hand variance and a correspondingly larger sample. Every interval reported in this paper is one standard error of the mean, $\bar{x} \pm \mathrm{SE}$ where $\mathrm{SE} = s/\sqrt{n}$; rows attributed to (Provost et al., 2026) are quoted as that paper reports them. GTO denotes game-theoretic optimal or approximate equilibrium play. A range is a distribution over possible private hands. Board texture describes how coordinated the public cards are. A blocker is a private card that reduces the opponent's combinations of a strong hand. A bluff-catcher is a hand that usually loses to value bets but can beat bluffs. SPR is stack-to-pot ratio, the effective remaining stack divided by the pot.

B A BRIEF HISTORY OF HUNL AND POKER AI

HUNL has served as a primary benchmark for imperfect-information game AI for over two decades. The evolution of HUNL agents illustrates a progression from interpretable but weak heuristics to powerful but opaque computational systems. PokerSkill represents a distinct point in this landscape: a training-free, solver-free agent that leverages rule-based knowledge through modern language models.

Era 1: Rule-based and expert systems (1990s–2005). The earliest poker AI programs used hand-crafted rules for betting decisions. The University of Alberta's poker research group developed Loki (Billings et al., 2002), which combined hand strength estimation, opponent modeling via neural networks, and rule-based betting logic. These systems were interpretable and required no large-scale computation, but their strategic depth was limited by the expressiveness of manually authored rules. They could not represent the complex conditional strategies needed for equilibrium play in no-limit variants.

Era 2: Abstraction and CFR (2005–2015). The introduction of counterfactual regret minimization (CFR) (Zinkevich et al., 2007) transformed poker AI. By iteratively minimizing regret over the game tree, CFR converges to Nash equilibrium in two-player zero-sum games. However, full HUNL has $\sim 10^{164}$ information sets (Johanson, 2013), far exceeding direct computation. Researchers developed action abstraction (Ganzfried & Sandholm, 2014) and card abstraction to reduce the game to tractable size, then solved the abstracted game with CFR variants (Lanctot et al., 2009; Brown & Sandholm, 2019a). Slumbot (Jackson, 2013) exemplifies this era: a strong HUNL agent built on large-scale CFR with abstraction, which won the 2018 Annual Computer Poker Competition.

Era 3: Depth-limited solving and real-time computation (2016–present). DeepStack (Moravčík et al., 2017) and Libratus (Brown & Sandholm, 2018) achieved superhuman HUNL play by combining CFR with depth-limited solving (Brown et al., 2018) and subgame re-solving (Brown & Sandholm, 2017). Rather than solving the entire game offline, these systems solve subgames in real time, using neural networks (DeepStack) or precomputed blueprints (Libratus) to estimate values beyond the solving depth. This approach requires millions of CPU core-hours for blueprint computation and significant real-time compute per decision. GTOWizard (Provost et al., 2026) represents the current state-of-the-art benchmark in this paradigm, beating Slumbot by 194 ± 41 mbb/hand. EVPA (Li

& Huang, 2025) and RL-CFR (Li et al., 2024) further improve abstraction and solving efficiency, enabling stronger equilibrium approximation with less computation.

Era 4: LLM-based agents (2023–present). The emergence of frontier language models opened a new direction: using LLMs as poker decision-makers without game-theoretic training. Early evaluations (Gupta, 2023; Zhuang et al., 2025) found that even GPT-4 deviates substantially from GTO play. ToolPoker (Lin et al., 2026) queries external solvers to improve LLM decisions but still depends on solver access. PokerSkill demonstrates that an LLM agent, equipped with structured rule-based skills but no solver or training, can achieve performance within the range of solver-based systems against a state-of-the-art benchmark.

Comparative positioning. Table 3 summarizes the key characteristics of representative HUNL agents across these eras.

Table 3: Comparison of HUNL poker AI approaches. "Training" refers to any offline computation (CFR iterations, RL, fine-tuning). "Interpretable" means the decision logic can be inspected by humans. The Slumbot row is GTOWizard's measurement of an earlier release of that bot and is not directly comparable to the PokerSkill row.

System	Training	Interpretable	Scales with LLM	vs. GTOWizard
Loki (rule-based)	None	Yes	N/A	$\gg$-200 mbb/h
Slumbot (CFR)	Millions CPU-hrs	No	No	-194 ± 41 mbb/h
DeepStack	Millions CPU-hrs	No	No	Near-human
Libratus	Millions CPU-hrs	No	No	Superhuman
GTOWizard	Large-scale	No	No	(reference)
GPT-5.5 (zero-shot)	None	Partial	Yes	-132 ± 25 mbb/h
PokerSkill (Claude Fable 5)	**None**	**Yes**	**Yes**	$\mathbf{-44 \pm 48}$ **mbb/h**

PokerSkill requires *zero training and zero solver access*. Because its decision quality depends on the underlying LLM, its performance moves with the base model: the five LLMs in Table 1 span -44 to -87 mbb/hand with identical rules.

C Prompt Layers

PokerSkill uses five prompt layers. P1 contains the game rules, output schema, and legal-action protocol. P2 contains preflop range guidance. P3 contains general postflop principles such as position, pot control, value/bluff separation, and bet sizing. P4 contains context-specific strategy keyed by board texture, action line, hand class, and role. P5 Blocker contains river bluff and bluff-catch guidance. The main paper describes the selection mechanism; this appendix records the implementation-facing taxonomy.

C.1 P1: System and Output Protocol

P1 is always included. It specifies that the model is playing HUNL, must choose only from legal actions, must respect the listed viable options, and must return a structured JSON action. The essential output schema is `{action, amount}`: `action` is one of fold, check, call, bet, raise, or all-in, and `amount` is a number or null. P1 also states that strategic reasoning should follow the computed ATTACK and DEFENSE budgets rather than inventing unconstrained bet sizes.

A minimal P1-style instruction is:

> You are a heads-up no-limit Texas hold'em agent. Use the situation analysis and selected poker skills below. Choose exactly one action from the legal and viable options. If betting or raising, output a legal amount. Return JSON only.

C.2 P2: Preflop Ranges

P2 is used only before the flop. The preflop scenario detector maps the action history to one of twelve scenario families such as unopened button, big blind versus open, 3-bet pot, 4-bet pot, limp

pot, or all-in confrontation. In the eleven covered families the agent samples its action from the frequency-annotated range for its starting hand and does not call the model; in the residual family the model receives the range text as compact advice for opening, calling, 3-betting, 4-betting, and folding. The model is only shown a chart in that residual scenario.

C.3 P3–P5 BLOCKER: POSTFLOP SKILLS

P3 contains stable postflop principles: separate value from bluffs, avoid betting marginal showdown hands without a reason, prefer position when realizing equity, use smaller bets on static boards, use protection bets on dynamic boards, and avoid illegal or meaningless bet sizes. P4 is retrieved by matching board texture, hero hand class, action-line scenario, position, and role. It gives local attack guidance, defense guidance, and warnings such as one-card flush danger or dominated kicker risk. P5 Blocker is included on the river and emphasizes that draws no longer have equity, blockers become more important, and marginal showdown hands should not be turned into value bets.

C.4 BOARD ANALYSIS AND SPECIAL BOARDS

Board classification covers texture, suit structure, and rank structure. Texture labels are `dry`, `slightly_wet`, `wet`, and `very_wet`; suit labels cover two-tone, monotone, flush-possible, and one-card-flush danger; rank labels cover paired, trips, double-paired, quads, straight-completing, and one-card straight threats. These labels set the attack/defense budget directly and determine which P4 warnings are retrieved.

Special boards are classified separately because ordinary hand ordering misleads there. On paired or trips boards a nominally strong hand can be dominated by full houses; on double-paired boards the kicker decides whether a high-card hand has showdown value; on quads or board-full-house textures private-card rank interaction matters more than the pair category; and on one-card flush or one-card straight boards a single private card completes a very strong hand, so medium showdown and high-card defense drop to zero. Each special board type receives a complete strategy override with its own pre-specified attack/defense budgets rather than an adjustment to the normal tables. A board such as $K\spadesuit 7\diamondsuit 2\clubsuit$ is dry and high-card dominated, while $9\clubsuit 8\clubsuit 6\diamondsuit$ is wet because it carries both flush and straight pressure.

Table 4: Board features computed by the context engine and their effect on the retrieved prompt.

Feature	Prompt effect
`dry`/`very_dry`	Small high-frequency c-bets, stable made-hand values, thinner value possible.
`wet`/`very_wet`	Larger sizes, lower c-bet frequency, protection and draw interaction emphasized.
Two-tone / 3-flush / 4-flush	Activates flush-draw, flush-complete, or one-card-flush warnings; non-flush made hands are downgraded.
Straight possible / one-card straight	Downgrades one-pair and two-pair hands; one-card straights distinguish top-end from low-end.
Paired board	Trips/full-house risk; paired-board penalties for pairs, limited bonus for high-card showdown.
Trips board	Quads/full-house/kicker override table replaces normal pair logic.
Double-paired / board full house	Private-card rank interaction decides whether hero has nuts, lower full house, shared board hand, or kicker-only showdown.
Board straight / board flush	Shared-board logic; private blocker or suited-card rank determines whether hero improves beyond the board.

C.5 PROMPT TEMPLATE

A simplified postflop prompt has this form:

> **Situation.** Street, pot, stacks, legal actions, hero cards, board, action history.
> **Computed context.** Pot type, position, initiative, SPR, board texture, hand class, draw class, action-line scenario, weighted pressure.

Budget verdict. Attack budget, defense budget, remaining budget, and reason.
Selected skills. P3 general principles; P4 entries for this board/hand/line; P5 river guidance if river.
Viable options. A short enumerated list of allowed strategic choices.
Output. Return JSON with one action and amount.

D ATTACK/DEFENSE BUDGET DETAILS

The budget system maps previous betting pressure to a scalar cost. Each hand class receives base attack and defense budgets, then pot type, board texture, position, SPR, and special-board modifiers adjust the budget. The action-line module compares the remaining budget to thresholds and emits a verdict plus viable actions. Close thresholds use fuzzy zones so the prompt can preserve both options rather than making discontinuous recommendations.

D.1 ATTACK AND DEFENSE BUDGET

Attack budget controls proactive betting and raising. Strong value hands receive high attack budget. Vulnerable but strong hands can retain attack budget on wet boards because protection has value. Marginal showdown and high-card showdown hands receive little or no attack budget because they should not bet for value. Draws can receive attack budget when fold equity and equity realization make bluffing reasonable, especially with strong draws or combo draws.

A reproduction should implement attack logic with three steps: assign a base attack budget from hand class, adjust by board danger, position, initiative, and SPR, and subtract weighted betting pressure already invested in the line. If remaining attack budget is clearly positive, betting or raising can appear in viable options. If it is near zero, only small bets or checks may remain. If it is negative, betting is removed unless an override applies for very strong hands or low-SPR commitment nodes.

Defense budget controls calls and defensive raises when facing bets. It starts from hand class and is modified by pot type, kicker, board danger, and position. Medium showdown hands defend more in limped and single-raised pots than in 3-bet and 4-bet pots. One-card flush or one-card straight boards can reduce defense to zero for high-card showdown hands. Paired boards can slightly increase showdown value, but this bonus does not turn dominated high-card hands into automatic calls.

D.2 BUDGET TABLES AND PROMPT TAXONOMY

Values are weighted-bet budgets before line-pressure subtraction. "Unlimited" is represented by a sentinel value in the implementation; ranges indicate kicker, rank, or texture-dependent variants. Board-feature labels and their effect on the prompt are listed in Section 3.1.

Complete hand-class ATT/DEF prompt content. The following is the exact prompt content provided to the LLM for each hand class. Each entry specifies the base ATT/DEF budgets and all context-dependent modifiers.

Hand class: nuts.

```
Full house+ / Nut flush / Nut straight (no-flush no-pair board) / Set (no-flush no-
    pair no-straight board):
  ATTACK: unlimited. DEFENSE: unlimited (always call/raise).
  Build pot aggressively. Consider overbet on nut advantage boards. Low SPR -> can
      slow-play one street to trap. RIVER: if you are last to act and have a raise
      option, ALWAYS raise -- calling with the nuts loses EV.
```

Hand class: flush.

```
Flush:
  [3-FLUSH BOARD (3 same suit)]:
    Nut flush: ATT/DEF unlimited.
```

```
    Big flush (high card >9): ATT 5 / DEF 6.
    Small flush (high card <=9): ATT 4 / DEF 5.
    [PAIRED BOARD]: Nut flush ATT 5 DEF 6. Big flush ATT 4.3 DEF 5.3. Small flush ATT
        3.5 DEF 4.5.
  [ONE-CARD FLUSH BOARD (4+ same suit)]:
    Nut: ATT/DEF unlimited. 2nd: ATT 4 DEF 5. 3rd: ATT 3 DEF 4.
    4th: ATT 2.4 DEF 3.5. 5th: ATT 2 DEF 3. 6-7th: ATT 1.5 DEF 2.5. 8-9th: ATT 1 DEF
        2.
    [PAIRED BOARD]: Nut flush ATT 4.5 DEF 5.5; others ATT/DEF -0.5.
```

Hand class: straight.

```
Straight:
  TWO-CARD STRAIGHT (both hole cards contribute):
    No flush: ATT 5.5 DEF 6.5. Flush possible (3-flush): ATT 3.5 DEF 4.5.
    4+ flush board: ATT 0 DEF 1.
    [PAIRED BOARD]: Nut straight + no flush -> ATT 5 DEF 6. Otherwise ATT/DEF -0.4.
  ONE-CARD STRAIGHT (1 hole card contributes):
    TOP-END (nut straight):
      No flush: NUTS. 3-flush: ATT 2.5 DEF 3.5. 4+flush: ATT 0 DEF 1.
    LOW-END (higher straight possible):
      No flush: ATT 2.5 DEF 3.5. 3-flush: ATT 1.5 DEF 2.5. 4+flush: ATT 0 DEF 0.5.
    [PAIRED BOARD]: ATT/DEF -0.6.
```

Hand class: set.

```
Set (pocket pair hits board):
  No flush, no straight possible: NUTS (ATT/DEF unlimited).
  No flush, has 2-card straight: 1 possibility ATT 5.5 DEF 6.5; 2 -> ATT 4.5 DEF 5.5;
      3+ -> ATT 4 DEF 5.
  3-flush board (no OCS): ATT 3.8 DEF 4.8, each additional straight possibility -0.3.
  One-card straight (OCS) 1 type: ATT 2.5 DEF 3.7. OCS 2+ types: ATT 1.5 DEF 2.7.
    OCS + 3-flush: ATT/DEF -0.8 from OCS base. OCS + 4+flush: ATT 0 DEF 0.5.
  4+ flush board (no OCS): ATT 0 DEF 1.
```

Hand class: trip.

```
Trips (board pair + one hole card):
  Dry board (no flush/OCS): ATT 4(2-kicker)-4.5(A-kicker) DEF 5.5, each 2-card str
      -0.3.
  3-flush board (no OCS): ATT 3-3.5 DEF 4.5, each 2-card str -0.3.
  OCS 1 type: ATT 2.1-2.4 DEF 3.5. OCS 2+ types: ATT 1.1-1.4 DEF 2.5.
    OCS + 3-flush: ATT/DEF -0.7 from OCS base. OCS + 4+flush: ATT 0 DEF 0.5.
  4+ flush board (no OCS): ATT 0 DEF 1.
  Kicker scales ATT proportionally: lowest(2)->base, highest(A)->max.
```

Hand class: two_pair.

```
Two pair (NO board pair):
  ATT/DEF by board texture (condition-matrix, highest-priority match):
  [4+ FLUSH (OCF) + ONE-CARD STRAIGHT]: ATT 0 / DEF 0.5.
  [4+ FLUSH (OCF), no OCS]: ATT 0 / DEF 1.
  [ONE-CARD STRAIGHT + 3-FLUSH]: 1 type OCS -> ATT 1.7 / DEF 2.8; 2+ types -> ATT 0.7
      / DEF 1.9.
  [ONE-CARD STRAIGHT, no flush]: 1 type OCS -> ATT 2.2 / DEF 3.3; 2+ types -> ATT 1.2
      / DEF 2.4.
```

```
  [3-FLUSH, no OCS]: rank-based ATT 2.7(r10)-3.6(r1) / DEF 3.7(r10)-4.7(r1); each 2-
      card straight -0.3.
  [DRY BOARD]: Rank 1 (top): ATT 5 / DEF 6.5. R2: 4.7/6. R3: 4.5/5.7. R4: 4.3/5.5.
    R5: 4.1/5.3. R6: 3.9/5.1. R7: 3.8/5. R8: 3.7/4.9. R9: 3.6/4.8. R10: 3.5/4.7.
    Each 2-card straight possibility -0.3 ATT/DEF.
  Top two pair on safe boards = go for stacks over all streets. Bottom two pair =
      still aggressive, but cautious on scary turn/river cards.
  NOTE: If board has a pair, your 'two pair' includes the shared board pair -- real
      strength ~ top pair/middle pair (see your classification above).
```

Hand class: overpair.

```
Overpair (pocket pair above all board cards):
  ATTACK / DEFENSE reference (by pot type and pair rank):
    SRP/Limp -- AA: ATT 3.5 / DEF 4.5. KK: ATT 3.4 / DEF 4.4. QQ: ATT 3.3 / DEF 4.3.
        JJ: ATT 3.2 / DEF 4.2. Others: ATT 3.1 / DEF 4.1.
    3BP+ -- AA: ATT 3.4 / DEF 4.5. KK: ATT 3.2 / DEF 4.3. QQ: ATT 3.0 / DEF 4.1. JJ:
        ATT 2.8 / DEF 3.8. TT: ATT 2.6 / DEF 3.7. Others: ATT 2.5 / DEF 3.5.
    4BP -- AA: ATT 3.4 / DEF 4.5. KK: ATT 3.1 / DEF 4.2. QQ: ATT 2.7 / DEF 3.7. JJ:
        ATT 2.4 / DEF 3.4. Others: ATT 2.1 / DEF 3.1.
  [SAFE BOARD -- no flush possible (rainbow/two-tone), no 1-card straight, no straight
      possible, not paired]:
    SRP: very strong -- triple barrel for value. In 3BP: triple barrel safe boards. In
        4BP: standard value, get stacks in.
  [PAIRED BOARD (board has a pair)]: Overpair on paired board: ATT/DEF -0.5. Still
      STRONG -- 2-3 streets of value. Flop paired board: c-bet and continue normally.
      Turn pairing: CHECK this specific street (opponent gains trips combos), but
      still bet the other streets. Overpair beats all one-pair hands and most bluffs
      -- do not play passively just because board is paired.
  [ONE-CARD FLUSH (board 4+ same suit, hero has NO flush)]: ATTACK 0, DEFENSE capped
      at 0.6. Near-worthless -- check/fold.
  [ONE-CARD STRAIGHT -- OPEN-ENDED]: drop 2.5 levels -> WEAK SHOWDOWN. Bet 1 street
      small at most, fold to heavy aggression.
  [ONE-CARD STRAIGHT -- GUTSHOT]: drop 1.5 levels -> THIN VALUE ONLY. Bet 1 street
      small, fold to heavy aggression.
  [FLUSH POSSIBLE BOARD (3+ same suit)]: drop 1.1(flop) / 0.9(turn) / 0.7(river)
      levels.
  [STRAIGHT POSSIBLE BOARD -- multiple combos]: drop 0.6(flop) / 0.5(turn) / 0.4(river
      ) levels.
  [STRAIGHT POSSIBLE BOARD -- only 1 combo]: drop 0.4(flop) / 0.3(turn) / 0.2(river)
      levels.
  NOTE: board texture penalties STACK, EXCEPT: [FLUSH POSSIBLE] and [ONE-CARD FLUSH]
      do NOT stack (use more severe); [STRAIGHT POSSIBLE] and [ONE-CARD STRAIGHT] do
      NOT stack (use more severe).
  Facing check-raise -> re-evaluate on all board types.
```

Hand class: top_pair.

```
Top pair:
  ATTACK / DEFENSE reference (by pot type and kicker):
    SRP -- TPTK: ATT 3 / DEF 4. TPSK: ATT 2.8 / DEF 3.8. 3rd kicker: ATT 2.6 / DEF
        3.6. 4th: ATT 2.4 / DEF 3.4. 5th: ATT 2.2 / DEF 3.2. Other: ATT 2.1 / DEF 3.1.
    3BP -- TPTK: ATT 2.9 / DEF 3.9. TPSK: ATT 2.6 / DEF 3.6. 3rd kicker: ATT 2.2 / DEF
        3.2. Other: ATT 1.9 / DEF 2.9.
    4BP+ -- TPTK: ATT 2.6 / DEF 3.6. TPSK: ATT 2.2 / DEF 3.2. 3rd kicker: ATT 1.8 /
        DEF 2.8. Other: ATT 1.6 / DEF 2.6.
  [SAFE BOARD -- no flush possible (rainbow/two-tone), no 1-card straight, no straight
      possible, not paired]:
    SRP TPGK / 3BP TPTK-TPSK: triple barrel for value on safe boards. Defend normal-
        size triple barrel.
    Other kickers: bet 2 streets, check 1 street for pot control.
```

```
  [PAIRED BOARD]: ATT/DEF -0.5. Opponent has more trips/full-house combos. Turn
      pairing: CHECK that specific street (opponent gains trips combos), bet the other
       streets.
  [ONE-CARD FLUSH (4+ same suit, hero has NO flush)]: drop 3.5 levels -> WEAK SHOWDOWN
      . Check, maybe call 1 small bet.
  [ONE-CARD STRAIGHT -- OPEN-ENDED]: drop 2.5 levels -> nearly WEAK SHOWDOWN. Check,
      maybe call 1 small bet.
  [ONE-CARD STRAIGHT -- GUTSHOT]: drop 1.5 levels -> THIN VALUE. Bet 1 street small at
       most.
  [FLUSH POSSIBLE BOARD (3+ same suit)]: drop 1.1(flop) / 0.9(turn) / 0.7(river)
      levels.
  [STRAIGHT POSSIBLE BOARD -- multiple combos]: drop 0.6(flop) / 0.5(turn) / 0.4(river
      ) levels.
  [STRAIGHT POSSIBLE BOARD -- only 1 combo]: drop 0.4(flop) / 0.3(turn) / 0.2(river)
      levels.
  NOTE: board texture penalties STACK, EXCEPT: [FLUSH POSSIBLE] and [ONE-CARD FLUSH]
      do NOT stack (use more severe); [STRAIGHT POSSIBLE] and [ONE-CARD STRAIGHT] do
      NOT stack (use more severe).
  Facing check-raise -> re-evaluate on all board types.
```

Hand class: second_pair.

```
Second pair (1 board overcard above hero's pair):
  ATTACK / DEFENSE reference (by pot type and kicker):
    SRP -- Pocket pair/top kicker: ATT 1.8 / DEF 2.8. 2nd kicker: ATT 1.7 / DEF 2.7. 3
        rd kicker: ATT 1.6 / DEF 2.6. Other: ATT 1.5 / DEF 2.5.
    3BP -- Pocket pair: ATT 1.3 / DEF 2.5. Top kicker: ATT 1.5 / DEF 2.3. 2nd kicker:
        ATT 1.3 / DEF 2.1. Other: ATT 1.2 / DEF 2.0.
    4BP+ -- Pocket pair: ATT 0 / DEF 2. Top kicker: ATT 1 / DEF 1.8. Other: ATT 0.9 /
        DEF 1.6.
  This hand has SOLID SHOWDOWN VALUE -- play for pot control.
  [PAIRED BOARD]: ATT/DEF -0.4.
  As aggressor: c-bet flop, bet 1 more street on safe boards. Max 2 streets of betting
      .
  River with second pair when double barreled before -> CHECK BACK.
  As defender SRP: check-call 2 streets. Facing <=40% pot bet -> call.
  In 3BP: check-call 1-2 streets, fold to continued heavy aggression.
  In 4BP: call once, fold to huge bets or continued aggression.
  [ONE-CARD FLUSH (4+ same suit, hero has NO flush)]: drop 3.5 levels -> Nearly TRASH.
  [ONE-CARD STRAIGHT -- OPEN-ENDED]: drop 2.5 levels.
  [FLUSH POSSIBLE BOARD (3+ same suit)]: drop 1.1(flop) / 0.9(turn) / 0.7(river)
      levels.
  [STRAIGHT POSSIBLE BOARD -- multiple combos]: drop 0.6(flop) / 0.5(turn) / 0.4(river
      ) levels.
  [STRAIGHT POSSIBLE BOARD -- only 1 combo]: drop 0.4(flop) / 0.3(turn) / 0.2(river)
      levels.
  NOTE: board texture penalties STACK, EXCEPT: [FLUSH POSSIBLE] and [ONE-CARD FLUSH]
      do NOT stack (use more severe); [STRAIGHT POSSIBLE] and [ONE-CARD STRAIGHT] do
      NOT stack (use more severe).
  Facing check-raise -> re-evaluate on all board types.
```

Hand class: third_pair.

```
Third pair (2 board overcards above hero's pair):
  ATTACK / DEFENSE reference (by pot type and kicker):
    SRP -- Top kicker/pocket pair: ATT 1.2 / DEF 2.2. Other: ATT 1.0 / DEF 2.0.
    3BP -- DEFENSE 1.5. 4BP+ -- DEFENSE 1.2.
    3BP+ -- Board-hit pair: ATTACK 1.5 (stab <=30% pot when opp shows weakness, check
        river). Pocket pair: ATTACK 0 DEFEND 0.6(3BP)/0.3(4BP+).
  3BP/4BP overcapped pocket pair (not hitting board, only 2 outs):
  [PAIRED BOARD]: ATT/DEF -0.3.
```

```
  This hand has MARGINAL SHOWDOWN VALUE.
  SRP: c-bet flop small (range bet ~33%), then CHECK remaining streets.
  3BP/4BP board-hit: can stab 1 street when opponent checks (showing weakness), then
      check river. Pocket pair: pure showdown.
  As defender SRP: facing c-bet -> call. Check-call 1-2 streets total. Fold to triple
      barrel.
  In 3BP/4BP: call once, fold to continued aggression.
  [ONE-CARD FLUSH (4+ same suit)]: drop 3.5 levels -> TRASH.
  [ONE-CARD STRAIGHT -- OPEN-ENDED]: drop 2.5 levels.
  [FLUSH POSSIBLE BOARD (3+ same suit)]: drop 1.1(flop) / 0.9(turn) / 0.7(river)
      levels.
  [STRAIGHT POSSIBLE BOARD -- multiple combos]: drop 0.6(flop) / 0.5(turn) / 0.4(river
      ) levels.
  [STRAIGHT POSSIBLE BOARD -- only 1 combo]: drop 0.4(flop) / 0.3(turn) / 0.2(river)
      levels.
  NOTE: board texture penalties STACK, EXCEPT: [FLUSH POSSIBLE] and [ONE-CARD FLUSH]
      do NOT stack (use more severe); [STRAIGHT POSSIBLE] and [ONE-CARD STRAIGHT] do
      NOT stack (use more severe).
```

Hand class: fourth_fifth_pair.

```
4th/5th pair (3-4 board overcards above hero's pair):
  ATTACK / DEFENSE reference (by pot type):
    SRP/Limp -- 4th pair: ATTACK 0.8 / DEFENSE 1.8. 5th pair: ATTACK 0.5 / DEFENSE
        1.5.
    3BP -- DEFENSE 1.0. 4BP+ -- DEFENSE 0.7.
    3BP+ -- Board-hit pair: ATTACK 1.5 (stab <=30% pot when opp shows weakness, check
        river). Pocket pair: ATTACK 0 DEFENSE 0.3(3BP)/0(4BP+).
  3BP/4BP overcapped pocket pair (not hitting board, only 2 outs):
  [PAIRED BOARD]: ATT/DEF -0.3.
  This hand has MINIMAL SHOWDOWN VALUE -- barely above trash.
  3BP/4BP board-hit: can stab 1 street when opponent checks, then check river.
  NEVER bet >=90% for value.
  As defender SRP: facing c-bet -> can call once. Fold to double barrel. Fold to
      triple barrel.
  In 3BP/4BP: fold to first significant bet (>=70% pot).
  [ONE-CARD FLUSH (4+ same suit)]: drop 3.5 levels -> TRASH.
  [ONE-CARD STRAIGHT -- OPEN-ENDED]: drop 2.5 levels -> TRASH.
  [FLUSH POSSIBLE BOARD (3+ same suit)]: drop 1.1(flop) / 0.9(turn) / 0.7(river)
      levels.
  [STRAIGHT POSSIBLE BOARD -- multiple combos]: drop 0.6(flop) / 0.5(turn) / 0.4(river
      ) levels.
  [STRAIGHT POSSIBLE BOARD -- only 1 combo]: drop 0.4(flop) / 0.3(turn) / 0.2(river)
      levels.
  NOTE: board texture penalties STACK, EXCEPT: [FLUSH POSSIBLE] and [ONE-CARD FLUSH]
      do NOT stack (use more severe); [STRAIGHT POSSIBLE] and [ONE-CARD STRAIGHT] do
      NOT stack (use more severe).
```

Hand class: nuts_high.

```
Nut high (A-high -- highest rank not on board):
  ATTACK 0. DEFENSE varies by pot type and kicker position (higher kicker -> more
      defense):
    Limp Pot: 0.8-1.2. SRP: 0.6-1.0. 3BP: 0.4-0.7. 4BP+: 0.1-0.4.
  Best non-pair showdown -- beats ALL bluffs and lower high cards.
  Do NOT bet for value. Check back or call small bets.
  PAIRED BOARD: DEFENSE boosted x1.35 (paired board = fewer pair combos for villain).
  [ONE-CARD FLUSH / ONE-CARD STRAIGHT]: DEFENSE -> 0 (too many hands beat you).
  [FLUSH POSSIBLE (3+ suit)]: penalty -1.0(flop) / -0.7(turn) / -0.4(river).
  [STRAIGHT POSSIBLE -- multi combos]: penalty -0.5(flop) / -0.4(turn) / -0.2(river).
  [STRAIGHT POSSIBLE -- 1 combo]: penalty -0.3(flop) / -0.2(turn) / -0.1(river).
```

```
  NOTE: board texture penalties STACK, EXCEPT: [FLUSH POSSIBLE] and [ONE-CARD FLUSH]
       do NOT stack (use more severe); [STRAIGHT POSSIBLE] and [ONE-CARD STRAIGHT] do
       NOT stack (use more severe).
```

Hand class: second_high.

```
Second high (K/Q-high -- 2nd-highest rank not on board, kicker T+ on unpaired board):
  ATTACK 0. DEFENSE varies by pot type and kicker position:
    Limp Pot: 0.4-0.7. SRP: 0.3-0.5. 3BP: 0.1-0.3. 4BP+: 0.
  Marginal showdown -- only beats bluffs. Do NOT bet.
  3BP+: fold to most bets.
  PAIRED BOARD: DEFENSE boosted x1.35.
  [ONE-CARD FLUSH / ONE-CARD STRAIGHT]: DEFENSE -> 0.
  [FLUSH POSSIBLE (3+ suit)]: penalty -1.0(flop) / -0.7(turn) / -0.4(river).
  [STRAIGHT POSSIBLE -- multi combos]: penalty -0.5(flop) / -0.4(turn) / -0.2(river).
  [STRAIGHT POSSIBLE -- 1 combo]: penalty -0.3(flop) / -0.2(turn) / -0.1(river).
  NOTE: board texture penalties STACK, EXCEPT: [FLUSH POSSIBLE] and [ONE-CARD FLUSH]
       do NOT stack (use more severe); [STRAIGHT POSSIBLE] and [ONE-CARD STRAIGHT] do
       NOT stack (use more severe).
```

Hand class: strong_draw.

```
Strong draw (combo draw on non-flushy / nut+ flush draw on flushy / flush draw rank>=J
     on non-flushy / OESD on rainbow non-straighty / combo draw flushy rank>=K):
  ATTACK: 4+ cumulative weighted bets (semi-bluff across multiple streets).
  DEFENSE (by opponent bet sizing, % of pot):
    Flop IP: defend up to 500% pot. Flop OOP: defend up to 400% pot.
    Turn IP: defend up to 190% pot. Turn OOP: defend up to 150% pot.
    Facing all-in: implied odds = ZERO. Need equity >= 60% pot odds to call.
    Facing check-raise: CALL (strong draws have enough equity to continue).
  As aggressor/normal: play aggressively, bet/raise as semi-bluff. Check-raise flop is
       default.
  As defender: check-call or check-raise depending on equity and position.
  Strategy similar across pot types. Exception: turn SPR <= 1.5 as IP -> check to
       preserve equity. COMBO RULE: pair + strong draw -> add ~2.0 extra defense to the
       PAIR/SHOWDOWN baseline. E.g. second pair (2.5) + strong draw (2.0) = defend
       ~4.5 cumulative bets.
```

Hand class: medium_strong_draw.

```
Medium-strong draw (flush draw rank<J on non-flushy / combo draw flushy rank J-Q /
     OESD on rainbow+straighty / OESD on two-tone non-straighty (2-card or clean
     conditions) / K+ flush draw on flushy):
  ATTACK: 3+ cumulative weighted bets (semi-bluff flop + selective turn barrel).
  DEFENSE (by opponent bet sizing, % of pot):
    Flop IP: defend up to 250% pot. Flop OOP: defend up to 200% pot.
    Turn IP: defend up to 100% pot. Turn OOP: defend up to 75% pot.
    Facing all-in: implied odds = ZERO. Need equity >= pot odds to call.
    Flop Facing check-raise: IP: defend up to 150% pot; OOP: defend up to 100% pot.
    Turn Facing check-raise: IP: defend up to 60% pot; OOP: defend up to 40% pot.
  As aggressor: semi-bluff flop, continue turn if equity holds or board improves.
  As normal: bet with high frequency as bluff. Check-raise flop is viable.
  In UNRAISED pots (no postflop raise): consider OOP check-raise as semi-bluff defense
       .
  COMBO RULE: pair + medium-strong draw -> add ~1.2 extra defense to the PAIR/SHOWDOWN
       baseline. E.g. second pair (2.5) + medium_strong_draw (1.2) = defend ~3.7
       cumulative bets.
```

Hand class: medium_draw.

```
Medium draw (decent flush draw rank T-Q on flushy board / OESD on two-tone+straighty /
     1-card OESD on two-tone non-straighty without clean conditions / combo draw
    flushy rank 8-T / non-flushy-board gutshot + overcards or backdoor flush on
    disconnected board):
  ATTACK: 1.5-3 cumulative weighted bets.
  DEFENSE (by opponent bet sizing, % of pot):
    Flop IP: defend up to 150% pot. Flop OOP: defend up to 120% pot.
    Turn IP: defend up to 60% pot. Turn OOP: defend up to 40% pot.
    Facing all-in: implied odds = ZERO. Need equity >= pot odds to call.
    Flop Facing check-raise: IP: defend up to 100% pot; OOP: defend up to 75% pot.
    Turn Facing check-raise: IP: defend up to 40% pot; OOP: defend up to 28% pot.
  As aggressor: double barrel flop+turn if equity holds.
  As normal: bet with high frequency as bluff.
  In UNRAISED pots (no postflop raise): consider OOP check-raise as semi-bluff defense
      .
  COMBO RULE: pair + medium draw -> add ~0.8 extra defense to the PAIR baseline. E.g.
     second pair (2.5) + medium draw (0.8) = defend ~3.3 cumulative bets. Do NOT fold
      pair+draw combos to single bets.
```

Hand class: medium_weak_draw.

```
Medium-weak draw (non-flushy 2-card gutshot / 1-card gutshot not at bottom of straight
     / moderate flush draw rank 6+ on flushy board / combo draw on flushy with high
    flush cards above / OESD+dual overcards on flushy / dual overcards + double
    backdoor):
  ATTACK: 1-2 cumulative weighted bets (bluff with equity)
  DEFENSE as % of pot by opponent bet sizing:
    Flop: IP ~94%, OOP ~78%
    Turn: IP ~40%, OOP ~26%
    Facing check-raise: Flop IP ~60%, OOP ~42%; Turn IP ~20%, OOP ~14%
  As aggressor: bet for value+protection on flop, check/fold turn if draw misses
  As defender: call flop bets profitably; turn calls require additional equity
  COMBO RULE: if paired (any pair + medium-weak draw), add ~0.7 extra defense budget
```

Hand class: weak_draw.

```
Weak draw (bottom gutshot on non-flushy / OESD on flushy without overcards / very
    small flush draw rank 2-5 on flushy / flushy gutshot + dual overcards /
    standalone double backdoor draws):
  ATTACK: 0.5-1 cumulative weighted bets (bluff only)
  DEFENSE as % of pot by opponent bet sizing:
    Flop: IP ~68%, OOP ~56%
    Turn: IP ~24%, OOP ~16%
    Facing check-raise: Flop IP ~40%, OOP ~28%; Turn: fold
  As defender: call small flop bets only; fold to turn aggression without improvement
  COMBO RULE: if paired (any pair + weak draw), add ~0.4 extra defense budget
```

Hand class: strong_overcard_draw.

```
Strong overcard draw (AK/AQ two premium overcards / KQ or AJ + backdoor straight / KJ/
    QJ/AT + 2-card backdoor flush or straight / overcard + real draw combo upgrade):
  ATTACK: 1 cumulative weighted bet (bluff with showdown equity)
  DEFENSE as % of pot by opponent bet sizing:
    Flop: IP ~80%, OOP ~65%
    Turn: IP ~35%, OOP ~25%
    Facing check-raise: Flop IP ~55%, OOP ~35%; Turn: fold
```

```
  As defender: defend flop at high frequency; turn defense requires improvement
      potential
  NOTE: texture downgrade may reduce this -- flushy board without flush draw or
      straighty board reduces tier
  3BP DOWNGRADE: all overcards drop 1 tier except AK, AQ/KQ/AJ+BD flush, AJ/KJ/AT+BD
      flush+BD straight
  4BP+ DOWNGRADE: all overcards drop 1 tier except AK+BD flush; many weaker overcards
      become trash
```

Hand class: medium_overcard_draw.

```
Medium overcard draw (naked dual overcards with rank sum > 19 like KQ/KJ/QJ/QT/JT/J9 /
      dual overcards + single backdoor / single overcard T+ with dual backdoor /
     single overcard T+ with single BD + another card above 2nd board rank / QJ on K-
     high / QT/JT/Q9 on K-high + BD flush / KQ/KJ on A-high without extras):
  ATTACK: 0.5-1 cumulative weighted bets (bluff only)
  DEFENSE as % of pot by opponent bet sizing:
    Flop: IP ~58%, OOP ~45%
    Turn: IP ~23%, OOP ~15%
    Facing check-raise: Flop IP ~28%, OOP ~18%; Turn: fold
  As defender: call small flop bets selectively; fold turn without improvement
  NOTE: texture downgrade may reduce this -- flushy board without flush draw or
      straighty board reduces tier
  3BP DOWNGRADE: drops to weak_overcard_draw (most fold to aggression)
  4BP+ DOWNGRADE: drops to weak_overcard_draw or trash depending on hand
```

Hand class: weak_overcard_draw.

```
Weak overcard draw (naked dual overcards with rank sum <= 19 like T8/T9/98 / dual
     overcards + single BD with sum <= 19 / single overcard T+ with BD but no between-
     pair / JT on Q-high / QT/JT/Q9 on K-high without BD / between-pair single
     overcard):
  ATTACK: 0.5 cumulative weighted bets (minimal bluff only)
  DEFENSE as % of pot by opponent bet sizing:
    Flop: IP ~35%, OOP ~25%
    Turn: IP ~15%, OOP ~9%
    Facing check-raise: fold
  As defender: fold to most aggression; only call very small bets on flop
  NOTE: texture downgrade may eliminate this entirely -- flushy or straighty board
      reduces to trash
  3BP DOWNGRADE: drops to trash (no overcard value in 3bp)
  4BP+ DOWNGRADE: drops to trash
```

Hand class: trash.

```
Pure trash (no pair, no overcard, no draw):
  ATTACK in flop/turn: 0-1 (and mostly small sizing stab/c-bet). DEFENSE: 0.
  Zero showdown value. Fold to any bet.
  IP STAB (aggressor checked to you): stab at HIGH frequency.
    Flop/turn: 20-30% pot. River: polarized bluff >60% pot (nothing to lose, max
        pressure).
  OOP: can make range bets (probe/c-bet at small sizing 20-30% pot). Can make >60% pot
       bet at river as bluff.
  RIVER BLUFF: If villain checks river to you, consider betting as bluff (zero
      showdown = nothing to lose). Villain's check signals weakness -- good bluff
      opportunity.
```

Pot-type adjustments. The following pot-type context is appended to the hand-class prompt to calibrate budget interpretation:

```
Limp pot: Very wide ranges. Value bet thinner (second pair can be strong). Bottom pair
    /A-high/K-high hands have more showdown value.

SRP: Standard ranges and thresholds. Top pair can be strong. Bottom pair/A-high can
    showdown or marginal bluff-catcher. Use standard sizing.

3-bet pot: Both players have NARROWER, STRONGER ranges. Recalibrate hand strength: top
     pair weak kicker ~ SRP second pair. Overpair below board top card is vulnerable.
     Two pair+ gains MORE value -- villain pays off with overpairs.

4BP+ pot: Very narrow ranges (premiums only). Range c-bet ~20% pot on most boards. Top
     pair ~ medium hand (only TPTK+ is strong). Overpair = bread-and-butter value
    hand. Low SPR -> stacks often committed, play straightforward.
```

River bluff and bluff-catch (P5). The following supplementary guidance is attached only on the river:

```
=== REFERENCE: River Bluff & Bluff-catch (lower priority -- use as supplementary
    guidelines) ===
[BLUFFING]
- Blocker: Bluff holding blockers to opponent's value hands (e.g., A spade blocks nut
    flush).
- No showdown value: Only bluff hands that lose at showdown.
- Consistent story: Bluff line must match strong hand narrative (bet-bet-bet).
- Bluff freq: bet_size/(pot+bet_size) of range should be bluffs.
- Good bluff: blocks opponent's value, doesn't block their bluffs or bluff-catchers.

[BLUFF-CATCHING]
- Need showdown value: Must beat opponent's bluff range (trash still folds).
- Prioritize calling stronger hands (kicker matters)
- Unblocker: Don't block opponent's bluffs (holding draw blockers = bad for catching).
- MDF: Defend >= pot/(pot+bet) to prevent opponent profiting with any bluff.
```

Bet-size to weighted pressure. Weighted pressure is computed from bet size as a percentage of the pot at the time of the bet, using a 47-step lookup table, abridged in Table 5: if the bet is below a threshold, the corresponding weight is returned.

Table 5: Weighted pressure $w(a_t)$ as a function of the bet size in percent of the pot at the time of the bet. The implementation uses a 47-step lookup table; representative thresholds are shown.

Threshold (% pot)	<5	<20	<32	<50	<67	<85	<100	<122	<150
Weight	0.04	0.30	0.50	0.70	0.85	1.00	1.10	1.25	1.40
Threshold (% pot)	<195	<300	<400	<500	<700	<1000	<1500	≥ 1500	
Weight	1.60	2.00	2.30	2.50	2.90	3.40	4.00	4.00	

Special-board override budgets.

Board	Override
Trips board	Quads/full-house are nuts; nut kicker ATT 0.5 DEF 1.5; second kicker DEF 0.8; lower kickers mostly trash.
Double-paired board	Higher-pair full house is nuts; lower full house ATT 2.5 DEF 3.5; flush/straight ATT 2 DEF 3; kicker-only tiers drop in 3BP/4BP.
Trips plus side cards	Quads nuts; matching side card or pocket pair forms full-house tiers; flush/straight downgraded below full house.
Quads board	Kicker decides: nuts high is nuts, second high ATT 1.5 DEF 2.5, third high ATT 0.5 DEF 1.5.
Full-house board	Most hands share board full house; only quads or higher full-house interaction plays aggressively.
Board flush	Highest private suited card sets rank: 2nd highest ATT 3 DEF 4 down to no suited card ATT 0 DEF 1.5.
Board straight	Usually chopped unless private card improves top end; flush context can dominate the shared straight.

E Evaluation Methodology and AIVAT Variance Reduction

E.1 AIVAT Overview

GTOWizard results are reported with AIVAT (Burch et al., 2018) in mbb/hand (see Appendix A for unit conversions).

AIVAT (Action-Informed Value Assessment Tool) is a variance-reduction technique for poker evaluation. Standard poker evaluation suffers from extremely high variance: a single hand can swing by hundreds of big blinds due to card runout, making raw win-rate estimates unreliable without tens of thousands of hands. AIVAT reduces this variance by using the evaluator's knowledge of its own strategy and of the chance distribution to construct a control variate.

AIVAT (Burch et al., 2018) is a control-variate estimator. For each hand it adds to the raw outcome v a sum of correction terms, one at every chance event and at every decision of the player whose strategy is known to the evaluator (here GTOWizard itself). At a node with information H where the acting distribution $p(\cdot \mid H)$ is known and $u(H, \cdot)$ is any value estimate fixed before the outcome is observed, the correction is

$$C(H, A) = \sum_a p(a \mid H)\, u(H, a) - u(H, A), \tag{2}$$

with A the realised card or action. Decisions of the evaluated agent, whose action distribution is unknown, receive no correction.

Proposition 3 (Unbiasedness of the corrected estimate)**.** *If $A \mid H \sim p(\cdot \mid H)$ at every corrected node, then $\mathbb{E}[C(H, A) \mid H] = 0$, so $\hat{v} = v + \sum C$ has the same expectation as v for any choice of the value estimates u. The estimates u affect only the variance; the identity requires the true acting distribution at each corrected node, which is why the agent's own decisions are left uncorrected.*

The value estimates are GTOWizard's counterfactual values under its reference strategy, which is why the reduction is large against a near-equilibrium opponent. GTOWizard reports a variance reduction of about $30\times$ relative to raw duplicate outcomes, roughly a $5.5\times$ reduction of the standard error; we did not re-measure it on our runs and use the evaluator's per-hand AIVAT values as returned. Standard errors are the sample standard deviation of per-hand AIVAT values divided by $\sqrt{n}$; contrasts between two runs use $\sqrt{\mathrm{SE}_1^2 + \mathrm{SE}_2^2}$, treating runs as independent. Anytime-valid variants stop an evaluation early while retaining the coverage guarantee (Li et al., 2026a); here every number is a fixed-sample AIVAT mean and standard error.

E.2 Opponents and Protocol

Agents. We evaluate five frontier LLMs with PokerSkill: GPT-5.5 with extended high reasoning (XHigh), Claude Opus 4.6 and Claude Opus 4.7 with maximum thinking, Claude Fable 5, and the open-weight DeepSeek-V4-Flash. We run two baselines. The default-prompt baseline gives each model the game state and a generic poker instruction, with no skill library and no budget system. The rule-based baseline runs the skill library with deterministic action selection, taking the first viable

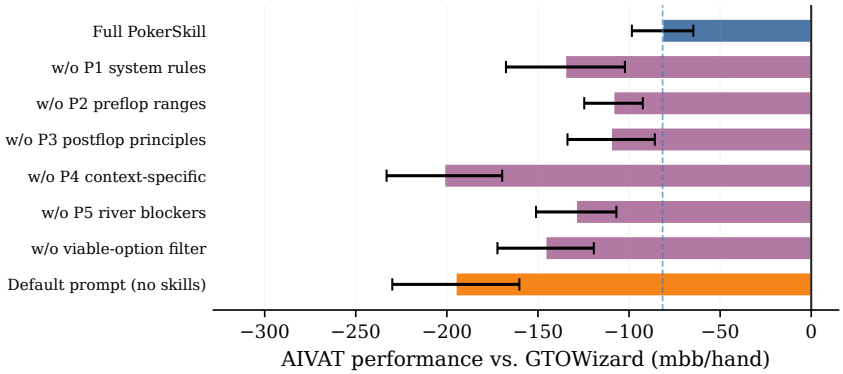


Figure 4: Leave-one-out ablation of the skill layers (DeepSeek-V4-Flash vs. GTOWizard, 10,000 AIVAT hands per bar, error bars = 1 SE), plotted from Table 2. The dashed line marks the full system.

action the budget system permits. Every agent returns a structured JSON action through forced tool use, and a validator rejects illegal or out-of-bounds sizes. No agent queries a solver during play. Experiments ran between April and September 2026 against the then-current API versions.

GTOWizard. GTOWizard is the only publicly available HUNL benchmark that provides AIVAT variance-reduced evaluation, and it represents current state-of-the-art GTO solving: it beats an earlier release of Slumbot by 194 ± 41 mbb/hand over 150,000 hands. Every model in Table 1 is evaluated against it with at least 5,000 AIVAT hands; the open-weight DeepSeek-V4-Flash is used for the 10,000-hand main run and the $6 \times 10{,}000$-hand ablation, whose default-prompt control has 5,000 hands.

Slumbot. Slumbot (Jackson, 2013) is the strongest publicly playable HUNL bot. We play the 2026 version hosted at `slumbot.com` through its public HTTP API with 200BB stacks. Slumbot offers no variance reduction, so outcomes are raw chip results and we report mbb/hand with $\pm$ one standard error over 400,000 hands for DeepSeek-V4-Flash with PokerSkill and 164,700 hands for the default-prompt control. Hands are dealt by the Slumbot server and the per-hand chip outcomes are pooled.

E.3 Experimental Details

LLM Model	Reasoning setting	Opponents
GPT-5.5 XHigh	Adaptive XHigh reasoning, temp 1.0	GTOWizard
Claude Opus 4.6	Max thinking (100K tokens), temp 1.0	GTOWizard
Claude Opus 4.7	Max thinking (100K tokens), temp 1.0	GTOWizard
Claude Fable 5	API default reasoning, temp 0.3	GTOWizard
DeepSeek-V4-Flash	API default reasoning, temp 0.3	GTOWizard, Slumbot

GPT-5.5 and the Claude Opus models use temperature 1.0 as their extended-thinking APIs require; Claude Fable 5 and DeepSeek-V4-Flash use the provider's default reasoning mode (no reasoning-effort or thinking-budget parameter is passed) with temperature 0.3. We control no random seeds: cards are dealt by the GTOWizard or Slumbot server, model sampling runs at the temperatures above, and preflop actions are sampled from the range table. Every configuration in Table 1 uses at least 5,000 hands; the DeepSeek-V4-Flash PokerSkill number pools two 5,000-hand sessions run on consecutive days with identical settings.

E.4 Ablation Estimand

Let μ_S be the expected AIVAT return of the agent built from the component set S against GTOWizard, and let $S^\star$ be the full set of six components (P1–P5 and the viable-option filter). Table 2 estimates the *conditional removal effects*

$$\Delta_j = \mu_{S^\star} - \mu_{S^\star \setminus \{j\}}, \qquad j \in S^\star, \tag{3}$$

each from an independent 10,000-hand run, so $\mathrm{SE}(\hat{\Delta}_j) = \sqrt{\mathrm{SE}^2_{S^\star} + \mathrm{SE}^2_{S^\star \setminus \{j\}}}$. A removal effect is not a standalone contribution: it measures what component j adds in the presence of the other five. If the components were additive, $\sum_j \Delta_j$ would equal $\mu_{S^\star} - \mu_\emptyset$, where $\mu_\emptyset$ is the default prompt. The observed excess is

$$\sum_j \hat{\Delta}_j - (\hat{\mu}_{S^\star} - \hat{\mu}_\emptyset) = 227 \pm 112 \text{ mbb/hand} \quad (z = 2.0), \tag{4}$$

where the standard error accounts for the six-fold reuse of the full-system run, $\mathrm{SE}^2 = 25\,\mathrm{SE}^2_{S^\star} + \sum_j \mathrm{SE}^2_{S^\star \setminus \{j\}} + \mathrm{SE}^2_\emptyset$. The estimated removal effects are 120 ± 36 (P4), 64 ± 31 (filter), 53 ± 37 (P1), 47 ± 28 (P5), 28 ± 29 (P3), and 27 ± 23 (P2); only the first two are more than two standard errors from zero, and the six comparisons are unadjusted.

F Additional Related Work

Retrieval-augmented and tool-augmented LLMs. RAG augments LLMs with external knowledge at inference time (Lewis et al., 2020), and Toolformer (Schick et al., 2023) teaches models to invoke tools. PokerSkill shares the principle of injecting task-relevant information at decision time, but its retrieval is *deterministic and rule-based*: the context engine computes exact game-state features and selects prompt fragments by logical condition.

Structured prompting. Chain-of-thought (Wei et al., 2022), Tree of Thoughts (Yao et al., 2023a), ReAct (Yao et al., 2023b), and Program-of-Thoughts (Chen et al., 2023) improve LLM performance through structured interfaces without parameter updates. These general-purpose methods help the model reason more carefully but do not tell it which poker principle applies to a given node; PokerSkill supplies that domain-specific selection. Broader evaluations confirm systematic LLM failures in game-theoretic settings (Huang et al., 2024; Hua et al., 2024).

G Complete Hand Prompt Traces

The following example is prompt traces from GPT-5.5 XHigh + PokerSkill. Each listing begins at the preflop decision and continues through the terminal decision, including the original prompt blocks, model responses, validated actions, and environment result lines.

```
============================================================
Street: preflop
Action: b 9.0

--- PROMPT ---
[RANGE TABLE] scenario=bb_vs_open hand=54s

--- GPT RESPONSE ---
{"action": "b", "amount": 9.0, "reasoning": "GTO range table"}

============================================================
Street: flop
Action: b 5.0

--- PROMPT ---
=== HUNL 200BB Hand ===

SITUATION:
- Street: FLOP
- Position: BB (OOP postflop)
- Pot type: 3BP (preflop aggressor: hero)
- You are the AGGRESSOR

YOUR HAND: 5c 4c (54s)
>>> Hand evaluation: STRONG DRAW <<<
```

```
BOARD: 7s 6h Jc
Board texture: RAINBOW board; MIXED high/low board

STACKS & POT:
- Pot: 18 | Total pot: 18
- Your stack: 191 | Villain: 191
- SPR: 10.6
- No bet to face

ACTION HISTORY:
  preflop: bet/raise to 2.75 -> bet/raise to 9 -> call

LEGAL ACTIONS: check (k), bet/raise (b) [1 - 191]

=== GENERAL PRINCIPLES (apply in priority order) ===

1. DEFENDER OOP -- CHECK OR DONK: If you are the DEFENDER and OOP, your DEFAULT action
      is CHECK to the aggressor (let them bet, then check-call or check-raise).
- DONK BET (15% pot) is viable in specific spots: flop low board (high card <= 8, no
    board pair), or turn/river when a non-top card pairs the board or board becomes
    double-paired.
- When viable options include DONK BET, weigh it against CHECK -- donk is not
    mandatory.

2. Flop C-BET: As AGGRESSOR, c-bet is default on most boards.
- High-card / paired boards: c-bet ~25% pot at high frequency.
- Default dry boards: c-bet ~65% pot at moderate frequency.
- Wet / low connected boards (e.g. T98, 753): OOP aggressor -> range check in SRP/3BP.
      IP aggressor -> lower freq with larger size.
- In 4BP+: range c-bet ~20% pot on almost all boards (your range is much stronger).

3. HAND STRENGTH FIRST: Check [YOUR HAND STRENGTH] below BEFORE applying MDF/pot odds.
- Pure trash (no pair, no draw) -> FOLD to any bet. MDF does NOT apply to unplayable
    hands.

4. POT CONTROL: Medium-strength hands control pot size.
- IP: check back for pot control (don't bloat pot with marginal hands). If the oop
    aggressor check to you, you can stab ~25% pot and then check to the river.
- OOP: check is always the mainly option. Sometimes if the IP aggressor check in the
    previous stage, you can probe/blocking bet ~25% pot for deny equity/river cheap
    showdown.
- Medium pairs, Small pairs, A-high -> Mostly aim for cheap showdown.

5. ATTACK BUDGET IS BINDING: The >>> ATTACK BUDGET <<< line in SITUATION ANALYSIS is
    the FINAL authority on whether to bet.
- When it says CHECK -> you MUST check, regardless of board texture, position, or
    scenario suggestions.
- When it says BET -> betting is allowed (choose sizing from options). Not mandatory.
- ATTACK BUDGET already integrates hand strength, board texture, and cumulative
    betting history into one decision.
- Do NOT override ATTACK BUDGET with your own judgment about hand strength or board
    safety.
- EXCEPTION -- RIVER BLUFF: Zero-showdown hands (trash/air) on RIVER may bluff when
    not facing a bet, regardless of budget. Nothing to lose by bluffing.
- EXCEPTION -- RIVER VALUE: On RIVER, if you are NOT the OOP player acting first (i.e
    ., IP in any case, or OOP facing a bet), value hands with ATTACK remaining >=1
    MUST bet or raise. Never flat-call with a value hand on the river unless you are
    OOP acting first (where check-raise trapping is allowed).
- EXCEPTION -- RIVER SHOWDOWN: On RIVER, if you are NOT OOP acting first, and ATTACK
    remaining <=0.5, and SPR >0.5: do NOT bet or raise -- check back (IP) or call/
    fold (facing bet). Thin value on the river risks a raise that costs far more than
     the thin value gained.

VALUE BETTING: Strong hands bet for value. Size by opponent's calling range.
```

```
V:B RATIO: River polarized range -- when betting 1x pot, value:bluff ~ 2:1. Adjust by
    bet size.
MDF: MDF = pot/(pot+bet). Defend >= this freq. But multi-street / OOP / range
    disadvantage -> can fold more.
POT ODDS: pot_odds = call/(pot+bet+call). Need equity > pot odds for profitable call.
IMPLIED ODDS: Deep stacks + drawing hands -> effective odds > pot odds. Set mining
    benefits.
COMMITMENT: SPR <4 -> committed with top pair+ (play straightforward, get stacks in).
    SPR >10 -> plan 2-3 streets of action.
BET SIZING:
- Min bet is at least 1BB. Bets can be in decimal form (rounded to 2 decimal places).
    In small pots (limp pot ~2BB), use 1BB as default bet (~50% pot).
- Flop range bet: ~33% pot (but never below 1BB)
- Turn/river polarized: ~75-100% pot
- Overbet (125-200% pot): polarized bet on nut advantage boards (turn/river only)
- Stab / blocking bet: ~20-35% pot
- MAX: never bet more than 200% of the pot postflop
- GEOMETRIC (PREFERRED): Prefer geometric sizing when it works out to <=150% pot per
    street. See [GEOMETRIC SIZING] below for exact amount. Use geometric as DEFAULT
    for value hands planning multi-street betting.
- RIVER ALL-IN: If SPR < 2 on river, consider all-in for both polarized value and
    bluffs.
- RIVER NO THIN VALUE: On river, do NOT bet < 50% pot for thin value (all-in excepted)
    . Either bet big (>=50% pot) for value, or check back.
- Facing smaller bets -> defend wider. Facing larger bets -> defend tighter.

PAIRED BOARD DEFENSE (flop only, OOP defender facing bet):
- Trips or non-overcard draw -> RAISE SMALL (33% pot). Do NOT flat-call.
- Pair / showdown hands -> normal CALL logic applies.

BET TYPE REFERENCE:
| Type | Size | Purpose |
| POLARIZED BET | 60-200% pot | VALUE or BLUFF |
| MEDIUM BET | 55-65% pot | THIN VALUE or BLUFF (river)|
| C-BET | 15-65% pot | DENY EQUITY (flop)|
| DELAY C-BET | 20-35% pot | DENY EQUITY (turn)|
| STAB | 20-35% pot | INFO + DENY |
| PROBE BET | 20-35% pot | INFO + DENY |
| BLOCK BET | 20-35% pot | BLOCK (river) |
| DONK BET | 15% pot | INFO (OOP defend) |
| SEMI-BLUFF LARGE | 67%+ pot | BLUFF (draw) |
| SEMI-BLUFF SMALL | 20-35% pot | BLUFF (draw) |

ATTACK/DEFENSE NOTATION:
- ATTACK N = you can profitably bet/raise a cumulative total of ~N weighted streets (
    larger bets count more than small bets).
- DEFENSE N = you can profitably call/continue against a cumulative total of ~N
    weighted bets from opponent. Defend at most N times, at minimum N-1 times.
- RAISE THRESHOLD: when ATTACK >= (number of postflop bets already placed + remaining
    streets + 1), consider raising when facing opponent's bet, then continue betting
    on non-dangerous runouts to river. Example: ATTACK>=4 hand on flop in BB facing a
     c-bet -> consider high-frequency check-raise (1 bet already placed + 2 remaining
     streets + 1 = 4).
- OOP CHECK-RAISE GUIDANCE (facing bet, non-river): OOP should CHECK-RAISE more than
    CALL when option available. Vulnerable hands (sets/two-pair on draw-heavy boards)
     -> HIGH frequency CR to deny equity. Invulnerable hands (dry boards) -> can mix,
     occasional trap OK.
- IP RAISE vs SMALL BET: IP facing small bets (<=50% pot) should RAISE more often --
    don't let opponent see cheap cards. IP facing check-raises or large bets as
    defender -> mostly CALL to keep ranges capped.
- LOW SPR RAISE: At low SPR, look at ATTACK remaining. Remaining enough -> raise to
    end hand quickly. Remaining VERY high at low SPR -> can slow-play (trap) since
    hand is too strong to need protection. If remaining <= 0, do NOT raise for value.
- GEOMETRIC SIZING PRIORITY: As polarized aggressor (NOT c-bet/stab), PREFER geometric
     sizing from [GEOMETRIC SIZING] section. Geometric gets stacks in efficiently
```

```
    across streets. Use geometric as DEFAULT for value hands; standard sizing (33%,
    50%, 75%) only when geometric unavailable or >150% pot.
- CHECK THRESHOLD: when ATTACK < (weighted bets so far + remaining streets + 1), and
    you can check (not facing a bet), consider CHECKING for pot control. Bet on later
     streets if opponent shows weakness. Example: ATTACK=2, turn weighted bets=1.35,
    remaining=1 -> threshold=3.35. ATTACK 2 < 3.35 -> check turn, bet river if
    villain checks.
- For DRAWS: defense is expressed as max bet sizing (% of pot) you can defend per
    street, because draw equity depends on pot odds vs bet size.
- COMBO (made hand + draw): when you have BOTH a pair/weak showdown AND a draw, your
    defense is stronger than either alone. Add ~1 (medium draw), ~0.7 (medium-weak
    draw), ~0.4 (weak draw), ~0.3 (strong overcard draw), ~0.2 (medium overcard draw)
    , ~0.1 (weak overcard draw) extra defense to the made-hand baseline. On PAIRED
    BOARD, use the tiered DEFENSE baseline (nut-high 1.5 / second-high 1.0 / third-
    high 0.8) for COMBO.
- DUAL CLASSIFICATION: when you have BOTH a made hand AND a draw, follow whichever
    classification gives HIGHER ATTACK/DEFENSE values.
- ALL-IN RULE: when facing all-in or calling means going all-in, implied odds = ZERO.
    Must have equity >= pot odds to call. Draws lose most value at all-in.
- POSITION ADJUSTMENT: On flop/turn, OOP made hands with SPR > 0.5 -> DEFENSE -0.3 (
    positional disadvantage costs equity across remaining streets).
- RIVER RE-RAISE PENALTY: On river, facing a raise after your bet (re-raise) ->
    DEFENSE -0.3 (opponent's raising range on river is extremely strong).

[YOUR HAND STRENGTH]
- Strong draw (combo draw on non-flushy / nut+ flush draw on flushy / flush draw rank
    >=J on non-flushy / OESD on rainbow non-straighty / combo draw flushy rank>=K):
  ATTACK: 4+ cumulative weighted bets (semi-bluff across multiple streets).
  DEFENSE (by opponent bet sizing, % of pot):
    Flop IP: defend up to 500% pot. Flop OOP: defend up to 400% pot.
    Turn IP: defend up to 190% pot. Turn OOP: defend up to 150% pot.
    Facing all-in: implied odds = ZERO. Need equity >= 60% pot odds to call.
    Facing check-raise: CALL (strong draws have enough equity to continue).
  As aggressor/normal: play aggressively, bet/raise as semi-bluff. Check-raise flop is
       default.
  As defender: check-call or check-raise depending on equity and position.
  Strategy similar across pot types. Exception: turn SPR <= 1.5 as IP -> check to
      preserve equity. COMBO RULE: pair + strong draw -> add ~2.0 extra defense to the
       PAIR/SHOWDOWN baseline. E.g. second pair (2.5) + strong draw (2.0) = defend
      ~4.5 cumulative bets.

=== SITUATION ANALYSIS (apply to current spot) ===
[OVERALL]
Before choosing your action, analyze the following about YOUR hand:
1. What is your absolute hand strength? (e.g., top pair good kicker, bottom pair,
    flush draw)
2. [Flop/Turn Only] What draws do you have? (flush draw, straight draw, backdoor draws
    , combo draws)
3. What is your relative hand strength? Relative strength depends on: absolute
    strength, board wetness (wet boards devalue non-nut hands), and opponent's action
     history (aggressive actions = stronger ranges)
4. What is your plan for future streets/opponent actions? (bet 3 streets, check one
    street, give up)
5. [Attention] READ the [YOUR HAND STRENGTH] section carefully -- it already
    classified your hand. Trust it over your own quick judgment.

[POT TYPE ADJUSTMENT]
- 3-bet pot: Both players have NARROWER, STRONGER ranges. Recalibrate hand strength:
    top pair weak kicker ~ SRP second pair. Overpair below board top card is
    vulnerable. Two pair+ gains MORE value -- villain pays off with overpairs.

[BOARD TEXTURE -- FLOP]
- Dry/rainbow flop: Aggressor c-bets high freq (~80%). Range advantage matters most.
    Small sizing (25-33% pot).
```

```
[ACTION LINE]
Scenario: F-A2 -- Aggressor OOP. We act first. -> C-bet or check (board-dependent).
Weighted postflop bets so far: 0 (hero: 0). Next bet counts as #1.
-> Hand strength baseline (3BP, 3bp -0.5, 1 weighted bets): second pair+ for value.
-> YOUR HAND: trash + strong draw Cumulative attacks so far: 0.0 weighted bets.
>>> ATTACK BUDGET: 5 (base 0.5, draw floor 5) | used: 0 | remaining: +5.0 | BET or
    CHECK-RAISE <<<
-> Viable options: C-BET (25% pot)

[BET SIZING REFERENCE (pot = 18.00 BB, use these BB amounts)]
- 20% pot (probe/block) = 3.60 BB
- 25% pot (small bet) = 4.50 BB
- 33% pot (1/3 pot) = 5.94 BB
- 50% pot (half pot) = 9.00 BB
- 66% pot (2/3 pot) = 11.88 BB
- 75% pot (3/4 pot) = 13.50 BB
- 100% pot (pot) = 18.00 BB
- 125% pot (1.25x pot) = 22.50 BB
- 150% pot (1.5x pot) = 27.00 BB
- 200% pot (2x pot) = 36.00 BB
- MAX ALL-IN (capped): 54.00 BB (300% pot cap, stack=191.00)

[GEOMETRIC SIZING]
- Geometric bet: 126% pot = 22.76 BB (SPR=10.6, 3 streets left). RECOMMENDED for multi
    -street value betting.

[POSITION]
- OOP leading: As aggressor/normal: lead with value + bluffs on favorable boards. As
    defender: lead only on boards where YOUR range improved more (rare).


What is your action? Respond in JSON: {"action": "X", "amount": N, "reasoning": "..."}

--- GPT RESPONSE ---
{"action":"b","amount":5,"reasoning":"Attack budget allows the listed 25% pot c-bet.
    With an open-ended straight draw as OOP 3-bet aggressor, small semi-bluff c-
    betting applies pressure while building equity realization."}

============================================================
Street: turn
Action: k

--- PROMPT ---
=== HUNL 200BB Hand ===

SITUATION:
- Street: TURN
- Position: BB (OOP postflop)
- Pot type: 3BP (preflop aggressor: hero)
- You are the AGGRESSOR

YOUR HAND: 5c 4c (54s)
>>> Hand evaluation: STRONG DRAW <<<
BOARD: 7s 6h Jc 6c
Board texture: TWO-TONE board (flush draw exists but flush NOT yet possible); PAIRED
    board (1 pair(s)); MIXED high/low board

STACKS & POT:
- Pot: 28 | Total pot: 28
- Your stack: 186 | Villain: 186
- SPR: 6.6
- No bet to face

ACTION HISTORY:
  preflop: bet/raise to 2.75 -> bet/raise to 9 -> call
```

```
  flop: bet/raise to 5 -> call

LEGAL ACTIONS: check (k), bet/raise (b) [1 - 186]

=== GENERAL PRINCIPLES (apply in priority order) ===

1. DEFENDER OOP -- CHECK OR DONK: If you are the DEFENDER and OOP, your DEFAULT action
     is CHECK to the aggressor (let them bet, then check-call or check-raise).
- DONK BET (15% pot) is viable in specific spots: flop low board (high card <= 8, no
    board pair), or turn/river when a non-top card pairs the board or board becomes
    double-paired.
- When viable options include DONK BET, weigh it against CHECK -- donk is not
    mandatory.

2. Flop C-BET: As AGGRESSOR, c-bet is default on most boards.
- High-card / paired boards: c-bet ~25% pot at high frequency.
- Default dry boards: c-bet ~65% pot at moderate frequency.
- Wet / low connected boards (e.g. T98, 753): OOP aggressor -> range check in SRP/3BP.
     IP aggressor -> lower freq with larger size.
- In 4BP+: range c-bet ~20% pot on almost all boards (your range is much stronger).

3. HAND STRENGTH FIRST: Check [YOUR HAND STRENGTH] below BEFORE applying MDF/pot odds.
- Pure trash (no pair, no draw) -> FOLD to any bet. MDF does NOT apply to unplayable
    hands.

4. POT CONTROL: Medium-strength hands control pot size.
- IP: check back for pot control (don't bloat pot with marginal hands). If the oop
    aggressor check to you, you can stab ~25% pot and then check to the river.
- OOP: check is always the mainly option. Sometimes if the IP aggressor check in the
    previous stage, you can probe/blocking bet ~25% pot for deny equity/river cheap
    showdown.
- Medium pairs, Small pairs, A-high -> Mostly aim for cheap showdown.

5. ATTACK BUDGET IS BINDING: The >>> ATTACK BUDGET <<< line in SITUATION ANALYSIS is
    the FINAL authority on whether to bet.
- When it says CHECK -> you MUST check, regardless of board texture, position, or
    scenario suggestions.
- When it says BET -> betting is allowed (choose sizing from options). Not mandatory.
- ATTACK BUDGET already integrates hand strength, board texture, and cumulative
    betting history into one decision.
- Do NOT override ATTACK BUDGET with your own judgment about hand strength or board
    safety.
- EXCEPTION -- RIVER BLUFF: Zero-showdown hands (trash/air) on RIVER may bluff when
    not facing a bet, regardless of budget. Nothing to lose by bluffing.
- EXCEPTION -- RIVER VALUE: On RIVER, if you are NOT the OOP player acting first (i.e
    ., IP in any case, or OOP facing a bet), value hands with ATTACK remaining >=1
    MUST bet or raise. Never flat-call with a value hand on the river unless you are
    OOP acting first (where check-raise trapping is allowed).
- EXCEPTION -- RIVER SHOWDOWN: On RIVER, if you are NOT OOP acting first, and ATTACK
    remaining <=0.5, and SPR >0.5: do NOT bet or raise -- check back (IP) or call/
    fold (facing bet). Thin value on the river risks a raise that costs far more than
     the thin value gained.

VALUE BETTING: Strong hands bet for value. Size by opponent's calling range.
V:B RATIO: River polarized range -- when betting 1x pot, value:bluff ~ 2:1. Adjust by
    bet size.
MDF: MDF = pot/(pot+bet). Defend >= this freq. But multi-street / OOP / range
    disadvantage -> can fold more.
POT ODDS: pot_odds = call/(pot+bet+call). Need equity > pot odds for profitable call.
IMPLIED ODDS: Deep stacks + drawing hands -> effective odds > pot odds. Set mining
    benefits.
COMMITMENT: SPR <4 -> committed with top pair+ (play straightforward, get stacks in).
    SPR >10 -> plan 2-3 streets of action.
BET SIZING:
```

```
- Min bet is at least 1BB. Bets can be in decimal form (rounded to 2 decimal places).
    In small pots (limp pot ~2BB), use 1BB as default bet (~50% pot).
- Flop range bet: ~33% pot (but never below 1BB)
- Turn/river polarized: ~75-100% pot
- Overbet (125-200% pot): polarized bet on nut advantage boards (turn/river only)
- Stab / blocking bet: ~20-35% pot
- MAX: never bet more than 200% of the pot postflop
- GEOMETRIC (PREFERRED): Prefer geometric sizing when it works out to <=150% pot per
    street. See [GEOMETRIC SIZING] below for exact amount. Use geometric as DEFAULT
    for value hands planning multi-street betting.
- RIVER ALL-IN: If SPR < 2 on river, consider all-in for both polarized value and
    bluffs.
- RIVER NO THIN VALUE: On river, do NOT bet < 50% pot for thin value (all-in excepted)
    . Either bet big (>=50% pot) for value, or check back.
- Facing smaller bets -> defend wider. Facing larger bets -> defend tighter.

PAIRED BOARD DEFENSE (flop only, OOP defender facing bet):
- Trips or non-overcard draw -> RAISE SMALL (33% pot). Do NOT flat-call.
- Pair / showdown hands -> normal CALL logic applies.

BET TYPE REFERENCE:
| Type | Size | Purpose |
| POLARIZED BET | 60-200% pot | VALUE or BLUFF |
| MEDIUM BET | 55-65% pot | THIN VALUE or BLUFF (river)|
| C-BET | 15-65% pot | DENY EQUITY (flop)|
| DELAY C-BET | 20-35% pot | DENY EQUITY (turn)|
| STAB | 20-35% pot | INFO + DENY |
| PROBE BET | 20-35% pot | INFO + DENY |
| BLOCK BET | 20-35% pot | BLOCK (river) |
| DONK BET | 15% pot | INFO (OOP defend) |
| SEMI-BLUFF LARGE | 67%+ pot | BLUFF (draw) |
| SEMI-BLUFF SMALL | 20-35% pot | BLUFF (draw) |

ATTACK/DEFENSE NOTATION:
- ATTACK N = you can profitably bet/raise a cumulative total of ~N weighted streets (
    larger bets count more than small bets).
- DEFENSE N = you can profitably call/continue against a cumulative total of ~N
    weighted bets from opponent. Defend at most N times, at minimum N-1 times.
- RAISE THRESHOLD: when ATTACK >= (number of postflop bets already placed + remaining
    streets + 1), consider raising when facing opponent's bet, then continue betting
    on non-dangerous runouts to river. Example: ATTACK>=4 hand on flop in BB facing a
     c-bet -> consider high-frequency check-raise (1 bet already placed + 2 remaining
     streets + 1 = 4).
- OOP CHECK-RAISE GUIDANCE (facing bet, non-river): OOP should CHECK-RAISE more than
    CALL when option available. Vulnerable hands (sets/two-pair on draw-heavy boards)
     -> HIGH frequency CR to deny equity. Invulnerable hands (dry boards) -> can mix,
     occasional trap OK.
- IP RAISE vs SMALL BET: IP facing small bets (<=50% pot) should RAISE more often --
    don't let opponent see cheap cards. IP facing check-raises or large bets as
    defender -> mostly CALL to keep ranges capped.
- LOW SPR RAISE: At low SPR, look at ATTACK remaining. Remaining enough -> raise to
    end hand quickly. Remaining VERY high at low SPR -> can slow-play (trap) since
    hand is too strong to need protection. If remaining <= 0, do NOT raise for value.
- GEOMETRIC SIZING PRIORITY: As polarized aggressor (NOT c-bet/stab), PREFER geometric
     sizing from [GEOMETRIC SIZING] section. Geometric gets stacks in efficiently
    across streets. Use geometric as DEFAULT for value hands; standard sizing (33%,
    50%, 75%) only when geometric unavailable or >150% pot.
- CHECK THRESHOLD: when ATTACK < (weighted bets so far + remaining streets + 1), and
    you can check (not facing a bet), consider CHECKING for pot control. Bet on later
     streets if opponent shows weakness. Example: ATTACK=2, turn weighted bets=1.35,
    remaining=1 -> threshold=3.35. ATTACK 2 < 3.35 -> check turn, bet river if
    villain checks.
- For DRAWS: defense is expressed as max bet sizing (% of pot) you can defend per
    street, because draw equity depends on pot odds vs bet size.
```

```
- COMBO (made hand + draw): when you have BOTH a pair/weak showdown AND a draw, your
    defense is stronger than either alone. Add ~1 (medium draw), ~0.7 (medium-weak
    draw), ~0.4 (weak draw), ~0.3 (strong overcard draw), ~0.2 (medium overcard draw)
    , ~0.1 (weak overcard draw) extra defense to the made-hand baseline. On PAIRED
    BOARD, use the tiered DEFENSE baseline (nut-high 1.5 / second-high 1.0 / third-
    high 0.8) for COMBO.
- DUAL CLASSIFICATION: when you have BOTH a made hand AND a draw, follow whichever
    classification gives HIGHER ATTACK/DEFENSE values.
- ALL-IN RULE: when facing all-in or calling means going all-in, implied odds = ZERO.
    Must have equity >= pot odds to call. Draws lose most value at all-in.
- POSITION ADJUSTMENT: On flop/turn, OOP made hands with SPR > 0.5 -> DEFENSE -0.3 (
    positional disadvantage costs equity across remaining streets).
- RIVER RE-RAISE PENALTY: On river, facing a raise after your bet (re-raise) ->
    DEFENSE -0.3 (opponent's raising range on river is extremely strong).

[YOUR HAND STRENGTH]
- Strong draw (combo draw on non-flushy / nut+ flush draw on flushy / flush draw rank
    >=J on non-flushy / OESD on rainbow non-straighty / combo draw flushy rank>=K):
  ATTACK: 4+ cumulative weighted bets (semi-bluff across multiple streets).
  DEFENSE (by opponent bet sizing, % of pot):
    Flop IP: defend up to 500% pot. Flop OOP: defend up to 400% pot.
    Turn IP: defend up to 190% pot. Turn OOP: defend up to 150% pot.
    Facing all-in: implied odds = ZERO. Need equity >= 60% pot odds to call.
    Facing check-raise: CALL (strong draws have enough equity to continue).
  As aggressor/normal: play aggressively, bet/raise as semi-bluff. Check-raise flop is
       default.
  As defender: check-call or check-raise depending on equity and position.
  Strategy similar across pot types. Exception: turn SPR <= 1.5 as IP -> check to
      preserve equity. COMBO RULE: pair + strong draw -> add ~2.0 extra defense to the
       PAIR/SHOWDOWN baseline. E.g. second pair (2.5) + strong draw (2.0) = defend
      ~4.5 cumulative bets.

=== SITUATION ANALYSIS (apply to current spot) ===
[OVERALL]
Before choosing your action, analyze the following about YOUR hand:
1. What is your absolute hand strength? (e.g., top pair good kicker, bottom pair,
    flush draw)
2. [Flop/Turn Only] What draws do you have? (flush draw, straight draw, backdoor draws
    , combo draws)
3. What is your relative hand strength? Relative strength depends on: absolute
    strength, board wetness (wet boards devalue non-nut hands), and opponent's action
     history (aggressive actions = stronger ranges)
4. What is your plan for future streets/opponent actions? (bet 3 streets, check one
    street, give up)
5. [Attention] READ the [YOUR HAND STRENGTH] section carefully -- it already
    classified your hand. Trust it over your own quick judgment.

[POT TYPE ADJUSTMENT]
- 3-bet pot: Both players have NARROWER, STRONGER ranges. Recalibrate hand strength:
    top pair weak kicker ~ SRP second pair. Overpair below board top card is
    vulnerable. Two pair+ gains MORE value -- villain pays off with overpairs.

[BOARD TEXTURE -- TURN]
- Turn pairs the board: As AGGRESSOR: reduce betting frequency on THIS street --
    especially when turn pairs a non-top flop card (check at higher frequency). But
    top pair+ still has 2-3 streets of total value -- if you only bet 1 street so far
    , you can still bet turn or river. Do NOT assume top pair is now worthless. Trips
    + can bet/raise for value. Check-raise from defender is credible.
- Neutral turn (only one flush draw possible, one suit has 2 cards): Only one flush
    draw is possible -- moderate texture. Made hands retain decent value. Can still
    bet for value with strong pairs. Check with medium hands if opponent could have
    the draw.

[ACTION LINE]
Scenario: T-A2 -- We bet flop OOP, villain called. -> Continue or check to defense.
```

```
Weighted postflop bets so far: 0.45 (hero: 0.45). Next bet counts as #1.45.
-> Hand strength baseline (3BP, 3bp -0.5, 1.45 weighted bets): second pair+ for value.
-> YOUR HAND: trash + strong draw Cumulative attacks so far: 0.5 weighted bets.
>>> ATTACK BUDGET: 5 (base 0.5, draw floor 5) | used: 0.45 | remaining: +4.5 | BET or
    CHECK-RAISE <<<
-> Viable options: CHECK

[BET SIZING REFERENCE (pot = 28.00 BB, use these BB amounts)]
- 20% pot (probe/block) = 5.60 BB
- 25% pot (small bet) = 7.00 BB
- 33% pot (1/3 pot) = 9.24 BB
- 50% pot (half pot) = 14.00 BB
- 66% pot (2/3 pot) = 18.48 BB
- 75% pot (3/4 pot) = 21.00 BB
- 100% pot (pot) = 28.00 BB
- 125% pot (1.25x pot) = 35.00 BB
- 150% pot (1.5x pot) = 42.00 BB
- 200% pot (2x pot) = 56.00 BB
- MAX ALL-IN (capped): 84.00 BB (300% pot cap, stack=186.00)

[GEOMETRIC SIZING]
- Geometric bet would be 176% pot (>150%), too large. Use standard sizing instead.

[POSITION]
- OOP leading: As aggressor/normal: lead with value + bluffs on favorable boards. As
    defender: lead only on boards where YOUR range improved more (rare).


What is your action? Respond in JSON: {"action": "X", "amount": N, "reasoning": "..."}

--- GPT RESPONSE ---
{"action":"k","reasoning":"Viable option is CHECK: the paired turn reduces OOP
    barreling frequency after villain called flop. With a strong draw, we can realize
     equity and potentially check-raise if villain bets."}

============================================================
Street: river
Action: b 21.0

--- PROMPT ---
=== HUNL 200BB Hand ===

SITUATION:
- Street: RIVER
- Position: BB (OOP postflop)
- Pot type: 3BP (preflop aggressor: hero)
- NEUTRAL -- aggressor checked prior street, no longer aggressive. You can probe, stab
    , or bet polarized based on hand strength

YOUR HAND: 5c 4c (54s)
>>> Hand evaluation: TRASH <<<
BOARD: 7s 6h Jc 6c 9h
Board texture: TWO-TONE board (flush draw exists but flush NOT yet possible); STRAIGHT
     POSSIBLE (2 two-card combos make straight); PAIRED board (1 pair(s)); MIXED high
    /low board

STACKS & POT:
- Pot: 28 | Total pot: 28
- Your stack: 186 | Villain: 186
- SPR: 6.6
- No bet to face

ACTION HISTORY:
  preflop: bet/raise to 2.75 -> bet/raise to 9 -> call
  flop: bet/raise to 5 -> call
```

```
  turn: check -> check

LEGAL ACTIONS: check (k), bet/raise (b) [1 - 186]

=== GENERAL PRINCIPLES (apply in priority order) ===

1. DEFENDER OOP -- CHECK OR DONK: If you are the DEFENDER and OOP, your DEFAULT action
     is CHECK to the aggressor (let them bet, then check-call or check-raise).
- DONK BET (15% pot) is viable in specific spots: flop low board (high card <= 8, no
    board pair), or turn/river when a non-top card pairs the board or board becomes
    double-paired.
- When viable options include DONK BET, weigh it against CHECK -- donk is not
    mandatory.

2. Flop C-BET: As AGGRESSOR, c-bet is default on most boards.
- High-card / paired boards: c-bet ~25% pot at high frequency.
- Default dry boards: c-bet ~65% pot at moderate frequency.
- Wet / low connected boards (e.g. T98, 753): OOP aggressor -> range check in SRP/3BP.
     IP aggressor -> lower freq with larger size.
- In 4BP+: range c-bet ~20% pot on almost all boards (your range is much stronger).

3. HAND STRENGTH FIRST: Check [YOUR HAND STRENGTH] below BEFORE applying MDF/pot odds.
- Pure trash (no pair, no draw) -> FOLD to any bet. MDF does NOT apply to unplayable
    hands.

4. POT CONTROL: Medium-strength hands control pot size.
- IP: check back for pot control (don't bloat pot with marginal hands). If the oop
    aggressor check to you, you can stab ~25% pot and then check to the river.
- OOP: check is always the mainly option. Sometimes if the IP aggressor check in the
    previous stage, you can probe/blocking bet ~25% pot for deny equity/river cheap
    showdown.
- Medium pairs, Small pairs, A-high -> Mostly aim for cheap showdown.

5. ATTACK BUDGET IS BINDING: The >>> ATTACK BUDGET <<< line in SITUATION ANALYSIS is
    the FINAL authority on whether to bet.
- When it says CHECK -> you MUST check, regardless of board texture, position, or
    scenario suggestions.
- When it says BET -> betting is allowed (choose sizing from options). Not mandatory.
- ATTACK BUDGET already integrates hand strength, board texture, and cumulative
    betting history into one decision.
- Do NOT override ATTACK BUDGET with your own judgment about hand strength or board
    safety.
- EXCEPTION -- RIVER BLUFF: Zero-showdown hands (trash/air) on RIVER may bluff when
    not facing a bet, regardless of budget. Nothing to lose by bluffing.
- EXCEPTION -- RIVER VALUE: On RIVER, if you are NOT the OOP player acting first (i.e
    ., IP in any case, or OOP facing a bet), value hands with ATTACK remaining >=1
    MUST bet or raise. Never flat-call with a value hand on the river unless you are
    OOP acting first (where check-raise trapping is allowed).
- EXCEPTION -- RIVER SHOWDOWN: On RIVER, if you are NOT OOP acting first, and ATTACK
    remaining <=0.5, and SPR >0.5: do NOT bet or raise -- check back (IP) or call/
    fold (facing bet). Thin value on the river risks a raise that costs far more than
     the thin value gained.

VALUE BETTING: Strong hands bet for value. Size by opponent's calling range.
V:B RATIO: River polarized range -- when betting 1x pot, value:bluff ~ 2:1. Adjust by
    bet size.
MDF: MDF = pot/(pot+bet). Defend >= this freq. But multi-street / OOP / range
    disadvantage -> can fold more.
POT ODDS: pot_odds = call/(pot+bet+call). Need equity > pot odds for profitable call.
IMPLIED ODDS: Deep stacks + drawing hands -> effective odds > pot odds. Set mining
    benefits.
COMMITMENT: SPR <4 -> committed with top pair+ (play straightforward, get stacks in).
    SPR >10 -> plan 2-3 streets of action.
BET SIZING:
```

```
- Min bet is at least 1BB. Bets can be in decimal form (rounded to 2 decimal places).
    In small pots (limp pot ~2BB), use 1BB as default bet (~50% pot).
- Flop range bet: ~33% pot (but never below 1BB)
- Turn/river polarized: ~75-100% pot
- Overbet (125-200% pot): polarized bet on nut advantage boards (turn/river only)
- Stab / blocking bet: ~20-35% pot
- MAX: never bet more than 200% of the pot postflop
- GEOMETRIC (PREFERRED): Prefer geometric sizing when it works out to <=150% pot per
    street. See [GEOMETRIC SIZING] below for exact amount. Use geometric as DEFAULT
    for value hands planning multi-street betting.
- RIVER ALL-IN: If SPR < 2 on river, consider all-in for both polarized value and
    bluffs.
- RIVER NO THIN VALUE: On river, do NOT bet < 50% pot for thin value (all-in excepted)
    . Either bet big (>=50% pot) for value, or check back.
- Facing smaller bets -> defend wider. Facing larger bets -> defend tighter.

PAIRED BOARD DEFENSE (flop only, OOP defender facing bet):
- Trips or non-overcard draw -> RAISE SMALL (33% pot). Do NOT flat-call.
- Pair / showdown hands -> normal CALL logic applies.

BET TYPE REFERENCE:
| Type | Size | Purpose |
| POLARIZED BET | 60-200% pot | VALUE or BLUFF |
| MEDIUM BET | 55-65% pot | THIN VALUE or BLUFF (river)|
| C-BET | 15-65% pot | DENY EQUITY (flop)|
| DELAY C-BET | 20-35% pot | DENY EQUITY (turn)|
| STAB | 20-35% pot | INFO + DENY |
| PROBE BET | 20-35% pot | INFO + DENY |
| BLOCK BET | 20-35% pot | BLOCK (river) |
| DONK BET | 15% pot | INFO (OOP defend) |
| SEMI-BLUFF LARGE | 67%+ pot | BLUFF (draw) |
| SEMI-BLUFF SMALL | 20-35% pot | BLUFF (draw) |

ATTACK/DEFENSE NOTATION:
- ATTACK N = you can profitably bet/raise a cumulative total of ~N weighted streets (
    larger bets count more than small bets).
- DEFENSE N = you can profitably call/continue against a cumulative total of ~N
    weighted bets from opponent. Defend at most N times, at minimum N-1 times.
- RAISE THRESHOLD: when ATTACK >= (number of postflop bets already placed + remaining
    streets + 1), consider raising when facing opponent's bet, then continue betting
    on non-dangerous runouts to river. Example: ATTACK>=4 hand on flop in BB facing a
     c-bet -> consider high-frequency check-raise (1 bet already placed + 2 remaining
     streets + 1 = 4).
- OOP CHECK-RAISE GUIDANCE (facing bet, non-river): OOP should CHECK-RAISE more than
    CALL when option available. Vulnerable hands (sets/two-pair on draw-heavy boards)
     -> HIGH frequency CR to deny equity. Invulnerable hands (dry boards) -> can mix,
     occasional trap OK.
- IP RAISE vs SMALL BET: IP facing small bets (<=50% pot) should RAISE more often --
    don't let opponent see cheap cards. IP facing check-raises or large bets as
    defender -> mostly CALL to keep ranges capped.
- LOW SPR RAISE: At low SPR, look at ATTACK remaining. Remaining enough -> raise to
    end hand quickly. Remaining VERY high at low SPR -> can slow-play (trap) since
    hand is too strong to need protection. If remaining <= 0, do NOT raise for value.
- GEOMETRIC SIZING PRIORITY: As polarized aggressor (NOT c-bet/stab), PREFER geometric
     sizing from [GEOMETRIC SIZING] section. Geometric gets stacks in efficiently
    across streets. Use geometric as DEFAULT for value hands; standard sizing (33%,
    50%, 75%) only when geometric unavailable or >150% pot.
- CHECK THRESHOLD: when ATTACK < (weighted bets so far + remaining streets + 1), and
    you can check (not facing a bet), consider CHECKING for pot control. Bet on later
     streets if opponent shows weakness. Example: ATTACK=2, turn weighted bets=1.35,
    remaining=1 -> threshold=3.35. ATTACK 2 < 3.35 -> check turn, bet river if
    villain checks.
- For DRAWS: defense is expressed as max bet sizing (% of pot) you can defend per
    street, because draw equity depends on pot odds vs bet size.
```

```
- COMBO (made hand + draw): when you have BOTH a pair/weak showdown AND a draw, your
    defense is stronger than either alone. Add ~1 (medium draw), ~0.7 (medium-weak
    draw), ~0.4 (weak draw), ~0.3 (strong overcard draw), ~0.2 (medium overcard draw)
    , ~0.1 (weak overcard draw) extra defense to the made-hand baseline. On PAIRED
    BOARD, use the tiered DEFENSE baseline (nut-high 1.5 / second-high 1.0 / third-
    high 0.8) for COMBO.
- DUAL CLASSIFICATION: when you have BOTH a made hand AND a draw, follow whichever
    classification gives HIGHER ATTACK/DEFENSE values.
- ALL-IN RULE: when facing all-in or calling means going all-in, implied odds = ZERO.
    Must have equity >= pot odds to call. Draws lose most value at all-in.
- POSITION ADJUSTMENT: On flop/turn, OOP made hands with SPR > 0.5 -> DEFENSE -0.3 (
    positional disadvantage costs equity across remaining streets).
- RIVER RE-RAISE PENALTY: On river, facing a raise after your bet (re-raise) ->
    DEFENSE -0.3 (opponent's raising range on river is extremely strong).

[YOUR HAND STRENGTH]
- Pure trash (no pair, no overcard, no draw):
  ATTACK in flop/turn: 0-1 (and mostly small sizing stab/c-bet). DEFENSE: 0.
  Zero showdown value. Fold to any bet.
  IP STAB (aggressor checked to you): stab at HIGH frequency.
    Flop/turn: 20-30% pot. River: polarized bluff >60% pot (nothing to lose, max
        pressure).
  OOP: can make range bets (probe/c-bet at small sizing 20-30% pot). Can make >60% pot
       bet at river as bluff.
  RIVER BLUFF: If villain checks river to you, consider betting as bluff (zero
      showdown = nothing to lose). Villain's check signals weakness -- good bluff
      opportunity.

=== SITUATION ANALYSIS (apply to current spot) ===
[OVERALL]
Before choosing your action, analyze the following about YOUR hand:
1. What is your absolute hand strength? (check [YOUR HAND STRENGTH] below -- trust its
     classification)
2. There are NO draws on the river -- all draws have either completed or missed.
3. What is your relative hand strength? Consider: board texture, opponent's action
    history across all streets, and what hands opponent is representing.
4. Is this a VALUE BET, BLUFF, or CHECK/CALL situation? Value bet strong hands, bluff
    with no showdown value + good blockers, check back medium hands.

[POT TYPE ADJUSTMENT]
- 3-bet pot: Both players have NARROWER, STRONGER ranges. Recalibrate hand strength:
    top pair weak kicker ~ SRP second pair. Overpair below board top card is
    vulnerable. Two pair+ gains MORE value -- villain pays off with overpairs.

[BOARD TEXTURE -- RIVER]
- Dry river (no flush possible): Board safe -- no draws completed. Made hand values
    stable. Bluff-catching viable with any showdown value. Bluffs need blockers to be
     credible.

[ACTION LINE]
Scenario: R-BC2 -- We bet flop OOP, checked turn. -> We can bet polarized or blocking
    bet or check.
Weighted postflop bets so far: 0.45 (hero: 0.45). Next bet counts as #1.45.
-> Hand strength baseline (3BP, 3bp -0.5, 1.45 weighted bets): second pair+ for value.
-> YOUR HAND: trash Cumulative attacks so far: 0.5 weighted bets.
>>> ATTACK BUDGET: 0.5 (base 0.5) | used: 0.45 | remaining: +0.1 | POLARIZED BET (
    bluff) or CHECK or BLOCK BET <<<
-> Viable options: POLARIZED BET (bluff) / BLOCK BET (20-35% pot) / CHECK

[BET SIZING REFERENCE (pot = 28.00 BB, use these BB amounts)]
- 66% pot (2/3 pot) = 18.48 BB
- 75% pot (3/4 pot) = 21.00 BB
- 100% pot (pot) = 28.00 BB
- 125% pot (1.25x pot) = 35.00 BB
- 150% pot (1.5x pot) = 42.00 BB
```

```
- 200% pot (2x pot) = 56.00 BB
- MAX ALL-IN (capped): 84.00 BB (300% pot cap, stack=186.00)

[GEOMETRIC SIZING]
- Geometric bet would be 664% pot (>150%), too large. Use standard sizing instead.

[POSITION]
- OOP leading: As aggressor/normal: lead with value + bluffs on favorable boards. As
    defender: lead only on boards where YOUR range improved more (rare).

=== REFERENCE: River Bluff & Bluff-catch (lower priority -- use as supplementary
    guidelines) ===
[BLUFFING]
- Blocker: Bluff holding blockers to opponent's value hands (e.g., As blocks nut flush
    ).
- No showdown value: Only bluff hands that lose at showdown.
- Consistent story: Bluff line must match strong hand narrative (bet-bet-bet).
- Bluff freq: bet_size/(pot+bet_size) of range should be bluffs.
- Good bluff: blocks opponent's value, doesn't block their bluffs or bluff-catchers.

[BLUFF-CATCHING]
- Need showdown value: Must beat opponent's bluff range (trash still folds).
- Prioritize calling stronger hands (kicker matters)
- Unblocker: Don't block opponent's bluffs (holding draw blockers = bad for catching).
- MDF: Defend >= pot/(pot+bet) to prevent opponent profiting with any bluff.


What is your action? Respond in JSON: {"action": "X", "amount": N, "reasoning": "..."}

--- GPT RESPONSE ---
{"action":"b","amount":21,"reasoning":"With zero showdown value, checking loses almost
     always, and villain's turn check caps many strong hands. Use a polarized 75% pot
     bluff rather than a small block bet to pressure pairs and missed floats."}
```