Beyond Outcome-Based Imperfect-Recall: Higher-Resolution Abstractions for Imperfect-Information Games

Yanchang Fu[*]
School of Artificial Intelligence
University of Chinese Academy of Sciences
Beijing, China
fuyanchang2020@ia.ac.cn

Qiyue Yin[†]
National Key Laboratory of Cognition and Decision Intelligence for Complex Systems
Institution of Automation, Chinese Academy of Sciences
Beijing, China
qyyin@nlpr.ia.ac.cn

Sengda Liu
National Key Laboratory of Cognition and Decision Intelligence for Complex Systems
Institution of Automation, Chinese Academy of Sciences
Beijing, China
shengda.liu@ia.ac.cn

Pei Xu
National Key Laboratory of Cognition and Decision Intelligence for Complex Systems
Institution of Automation, Chinese Academy of Sciences
Beijing, China
pei.xu@ia.ac.cn

Kaiqi Huang[‡]
National Key Laboratory of Cognition and Decision Intelligence for Complex Systems
Institution of Automation, Chinese Academy of Sciences
Beijing, China
kaiqi.huang@nlpr.ia.ac.cn

ABSTRACT

Hand abstraction is crucial for scaling imperfect-information games (IIGs) such as Texas Hold'em, yet progress is limited by the lack of a formal task model and by evaluations that require resource-intensive strategy solving. We introduce **signal observation ordered games (SOOGs)**, a subclass of IIGs tailored to hold'em-style games that cleanly separates signal from player action sequences, providing a precise mathematical foundation for hand abstraction. Within this framework, we define a **resolution bound**-an information-theoretic upper bound on achievable performance under a given abstraction algorithm. Using the bound, we show that mainstream outcome-based imperfect-recall algorithms suffer substantial losses by arbitrarily discarding historical information; we formalize this behavior via **potential-aware outcome Isomorphism (PAOI)** and prove that PAOI characterizes their resolution bound. To overcome this limitation, we propose **full-recall outcome isomorphism (FROI)**, which integrates historical information to raise the bound and improve policy quality. Experiments on a hold'em game benchmark confirm that FROI consistently outperforms outcome-based imperfect-recall baselines. Our research provides practical guidance for further designing higher-resolution abstraction algorithms in IIGs.

KEYWORDS

Game Theory; Imperfect-Information Games; Hand Abstraction

ACM Reference Format:
Yanchang Fu, Qiyue Yin, Sengda Liu, Pei Xu, and Kaiqi Huang. 2026. Beyond Outcome-Based Imperfect-Recall: Higher-Resolution Abstractions for Imperfect-Information Games. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 9 pages. https://doi.org/10.65109/AKDO5185

1 INTRODUCTION

In recent years, artificial intelligence (AI) for **heads-up no-limit hold'em** (HUNL)—a two-player poker variant and classic testbed for **imperfect-information games** (IIGs, where players lack full knowledge of the current game state)—has reached groundbreaking milestones. AI systems such as DeepStack, Libratus, and Pluribus [3, 5, 18] have defeated top human professionals, demonstrating AI's potential in mastering complex strategic reasoning.

Hand abstraction has become a key game-simplification technique for building AI for large-scale IIGs, especially HUNL. This is because state-of-the-art (SOTA) IIG solvers are based on **counterfactual regret minimization (CFR)** [27] and its variants [4, 16, 23, 26], whose exorbitant spatial overhead surpasses current computational resources. Among existing hand abstraction methods, **potential-aware abstraction with earth mover's distance (PAAEMD)** [6] is the SOTA and a critical component of high-performance HUNL AIs like DeepStack, Libratus, and Pluribus.

However, hand abstraction—an essentially engineering-driven technique—suffers from three key limitations. First, it lacks a formal mathematical foundation: no prototype in established IIG frameworks can characterize its game-simplification logic, leaving no systematic tool to guide the design of hand abstraction algorithms (with development relying heavily on empirical trial and error). Second, its evaluation is indirect and computation-intensive: a strategy

---

[*] Also with , National Key Laboratory of Cognition and Decision Intelligence for Complex Systems, Institution of Automation, Chinese Academy of Sciences.
[†] Corresponding author
[‡] Also with , School of Artificial Intelligence, University of Chinese Academy of Sciences.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/AKDO5185

must first be solved for the abstracted game to be assessed in the original game. This undermines its core goal of simplifying strategy solving—even solving the abstracted game incurs substantial computational overhead, which remains a prerequisite for evaluation nonetheless. Third, mainstream abstraction algorithms (e.g., PAAEMD) fully discard historical game information, leading to severe information loss. This inherent flaw restricts performance, as critical prior-state context (vital for distinguishing strategically distinct hands) is omitted, impairing the abstraction's ability to preserve the original game's structure.

To address these issues, this paper presents three key contributions, structured as follows.

First, to tackle the lack of refined modeling for hand abstraction, we construct **signal observation ordered games (SOOGs)**—a subset of IIGs tailored to capture the core characteristics of Texas Hold'em-style games (referred to as **hold'em games** throughout, including HUNL). Based on SOOGs, we further develop a formal mathematical model for hand abstraction (termed **signal observation abstraction**), detailed in Section 3 following essential background and notations in Section 2.

Second, to resolve the inefficiency of indirect evaluation for hand abstraction algorithms, we introduce **resolution bound**—a indicator that directly quantifies their maximum performance potential—and verify its validity as an evaluation metric in Section 4.

Third, we design two targeted hand abstractions: **potential-aware outcome isomorphism (PAOI)** and **full-recall outcome isomorphism (FROI)**. PAOI illustrates flaws in mainstream algorithms and is proven to be their resolution bound, explicitly revealing critical limitations (notably arbitrary historical information discard). FROI, by contrast, demonstrates how integrating historical information mitigates these flaws and improves abstraction quality. Section 4 details implementations and analyses of both methods.

We conclude that historical information is critical to hand abstraction performance. Section 5 provides experimental validation, confirming FROI's significant superiority over PAOI and underscoring its value for advancing hand abstraction techniques.

1.1 Related Works

Hand abstraction for hold'em games originates from the foundational work of Shi and Littman [20] and Billings et al. [1]. Gilpin and Sandholm [7] pioneered automatic methods in this field, shifting from manual hand categorization. Later, Gilpin and Sandholm [9] introduced **lossless isomorphism (LI)** for hand abstraction—constructing equivalence classes via card suit rotation. However, LI's low compression rate redirected research to lossy methods based on hand outcome win rates and **imperfect-recall** [25], including expected hand strength (EHS) [8], potential-aware abstractions (PAAs) [10], and their SOTA variant PAAEMD [6]. Notably, these lossy algorithms all discard historical information arbitrarily, relying on extreme imperfect-recall—a critical limitation we address in this work.

For IIG modeling relevant to hold'em games, Gilpin and Sandholm [9]'s **games with ordered signals** framework is the closest, yet it suffers from cumbersome phase modeling, limited generalizability, and no formal hand abstraction models. Subsequent IIG abstraction research (e.g., Waugh et al. [24]; Kroer and Sandholm [13, 14]) focused on abstraction but conflated it with techniques like action abstraction. Works on public states (Johanson et al. [11]) and public belief states (Šustr et al. [22]; Lisý and Bowling [17]) first noted the uniqueness of public action sequences in hold'em games (studied independently), while recent factored-observation stochastic games (FOSG) models [12, 19] further recognized independent observation modeling—applied to scenarios beyond hold'em games. But these modeling works did not involve hand abstraction.

2 BACKGROUND

In this section, we first provide an overview of hold'em games, then introduce the core concept of imperfect-information games (IIGs), and finally touch on content related to strategy solving for them.

2.1 Hold'em Games

Hold'em games are poker variants where players combine private (hole) and community (board) cards to form the strongest x-card hand. With multiple betting phases, they require sequential decisions based on personal hands and opponent action observations, with outcomes decided by all but one player folding or a showdown of remaining players' hands.

Heads-up limit hold'em (HULH) and HUNL are key IIG benchmarks for AI research—sharing core card-dealing and hand-formation rules, differing only in action spaces (HUNL's flexible betting expands its action space and strategic complexity). Focused on hand abstraction, this paper only involves HULH (rules in Appendix A.1)[1]. Other Hold'em variants serve as AI toy games, including Kuhn poker [15], Leduc Hold'em [21], and Flop Hold'em [2]; These simplified versions keep core strategy features and lower complexity, supporting practical AI experiments and analysis.

2.2 Imperfect-Information Games

Imperfect-information games are a standard mathematical model for modeling hold'em games.

Definition 2.1 (Imperfect-Information Game (IIG)). An imperfect-information game $\mathcal{G}$ is described by the following components:

- $\mathcal{N}_c = \mathcal{N} \cup \{c\}$: A finite set of players. $\mathcal{N} = \{1, \dots, N\}$ denotes rational players (actual game participants), and the special player c represents random factors (referred to as **nature** or **chance**).
- $\mathcal{A}$: The set of all possible actions available to players in $\mathcal{N}_c$ throughout the game.
- H: A finite set of histories, where each history is a sequence of successive player actions. The empty sequence $h^o \in H$ is the initial history. $h \sqsubset h'$ denotes h as a predecessor of h' (and h' its successor)—formally, there exists $a_1, a_2, \dots, a_k \in \mathcal{A}$ such that $h' = h \cdot a_1 \cdot a_2 \cdot \dots \cdot a_k$ (i.e., h' is h with the actions appended).
- $Z \subseteq H$: The set of terminal histories (histories with no successors), where the game terminates.
- $\rho : H \setminus Z \mapsto \mathcal{N}_c$: An action assignment function that specifies exactly one player to act at each non-terminal history. It induces a partition $\{H_1, \dots, H_N, H_c\}$ of $H \setminus Z$, where $H_i = \{h \in H \setminus Z \mid \rho(h) = i \in \mathcal{N}_c\}$ (the set of histories where player i acts).
- $A : H \setminus Z \mapsto 2^{\mathcal{A}}$: A legal action function that maps each non-terminal history to its set of legal actions. Here, $2^{\mathcal{A}}$ denotes the power set of $\mathcal{A}$ (i.e., all subsets of $\mathcal{A}$).

---

[1] A full version with appendices is available at https://arxiv.org/pdf/2510.15094.

- ζ: A chance probability function. For each non-terminal history $h \in H_c$ (where nature acts), $\zeta(h,\cdot) : A(h) \to [0,1]$ defines a probability distribution over $A(h)$, satisfying the normalization condition $\sum_{a \in A(h)} \zeta(h,a) = 1$.
- $u = (u_i)_{i \in \mathcal{N}}$: A tuple of utility functions. For each player $i \in \mathcal{N}$, $u_i : Z \to \mathbb{R}$ assigns a real-valued payoff to i at every terminal history $z \in Z$ (determining i's reward when the game ends at z).
- $\mathcal{I} = (\mathcal{I}_i)_{i \in \mathcal{N}}$: A collection of information partitions. For each $i \in \mathcal{N}$, $\mathcal{I}_i$ partitions H_i into disjoint **information sets (infosets)**. If $h, h' \in H_i$ belong to the same infoset $I \in \mathcal{I}_i$, rational player i cannot distinguish h and h'; this implies $A(h) = A(h')$, so we overload notation as $A(I) := A(h) = A(h')$.

2.3 Strategies and Solution Concepts for IIGs

In an IIG, each player $i \in \mathcal{N}$ selects a **(behavioral) strategy** σ_i, where Σ_i denotes the set of all such strategies for i. A strategy σ_i describes the probability distribution over actions chosen at each infoset: formally, $\sigma_i = \{\sigma_i(I,\cdot) \mid I \in \mathcal{I}_i\}$, where $\sigma_i(I,\cdot) : A(I) \to [0,1]$ satisfies $\sum_{a \in A(I)} \sigma_i(I,a) = 1$ (normalization) for each infoset $I \in \mathcal{I}_i$. When all rational players select their strategies, they form a strategy profile $\sigma = (\sigma_1, \dots, \sigma_N)$, with the set of all possible profiles denoted by $\Sigma = \times_{i \in \mathcal{N}} \Sigma_i$.

To compute payoffs under a strategy profile $\sigma \in \Sigma$, we first define the reaching probabilities of histories. For any history $h \in H$, let $\pi_c(h) = \prod_{h' \sqsubset h, h' \cdot a \sqsubseteq h, \rho(h') = c} \zeta(h', a)$ denote the contribution of nature to the probability of reaching h (where $h' \sqsubseteq h$ extends $h' \sqsubset h$ to include the case $h' = h$); for a player $i \in \mathcal{N}$, let $\pi_i^\sigma(h) = \prod_{h' \sqsubset h, h' \cdot a \sqsubseteq h, \rho(h') = i} \sigma_i(I[h'], a)$ denote the contribution of i's strategy σ_i to reaching h (with $I[h] \in \mathcal{I}_i$ being the infoset containing h). Since nature's choices are independent of players' strategies, $\pi_c(h)$ is invariant to σ (and thus we omit σ in its notation, avoiding distinction between $\pi_c^\sigma(h)$ and $\pi_c(h)$). The total probability of reaching h is then $\pi^\sigma(h) = \prod_{i \in \mathcal{N}_c} \pi_i^\sigma(h)$, and the expected payoff of player i under σ is $u_i(\sigma) = \sum_{z \in Z} \pi^\sigma(z) \cdot u_i(z)$.

An IIG has **perfect-recall** if every player $i \in \mathcal{N}$ never forgets past actions or observations: for any $h, h' \in I \in \mathcal{I}_i$, every predecessor $h'' \sqsubset h$ with $h'' \in H_i$ has a unique counterpart $h''' \sqsubset h'$ in the same infoset (i.e., exists $I' \in \mathcal{I}_i$ s.t., $h'', h''' \in I'$), and the actions taken by player i at h'' and h''' (to reach h and h') are identical. Under perfect-recall, a strategy profile $\sigma^* = (\sigma_1^*, \dots, \sigma_N^*)$ is a **Nash equilibrium** if for all $i \in \mathcal{N}$ and $\sigma_i \in \Sigma_i$, $u_i(\sigma^*) \geq u_i(\sigma_i, \sigma_{-i}^*)$ (where σ_{-i}^* denotes the strategy profile of all players except i). Every finite perfect-recall IIG guarantees at least one such equilibrium.

For 2-player zero-sum IIGs ($N = 2$, $u_1(z) = -u_2(z)$ for all $z \in Z$) with perfect-recall, the **Minimax Theorem** holds: there exists a unique game value v, and strategies $\sigma_1^* \in \Sigma_1$ and $\sigma_2^* \in \Sigma_2$ such that

$$\max_{\sigma_1 \in \Sigma_1} \min_{\sigma_2 \in \Sigma_2} u_1(\sigma_1, \sigma_2) = \min_{\sigma_2 \in \Sigma_2} \max_{\sigma_1 \in \Sigma_1} u_1(\sigma_1, \sigma_2) = v.$$

While multiple Nash equilibria may exist, all yield this unique value v, and any pair of maximin and minimax strategies (σ_1^*, σ_2^*) forms a Nash equilibrium. **Exploitability** measures how far a strategy profile $\sigma = (\sigma_1, \sigma_2)$ is from any Nash equilibrium $\sigma^* = (\sigma_1^*, \sigma_2^*)$, quantified by the maximum gains each player can achieve by deviating from σ). Its exploitability is defined as:

$$\epsilon(\sigma) = \frac{\overbrace{\left(u_1(\sigma^*) - \min_{\sigma_2' \in \Sigma_2} u_1(\sigma_1, \sigma_2')\right)}^{\epsilon_1(\sigma)} + \overbrace{\left(u_2(\sigma^*) - \min_{\sigma_1' \in \Sigma_1} u_2(\sigma_1', \sigma_2)\right)}^{\epsilon_2(\sigma)}}{2},$$

where a profile with $\epsilon(\sigma) \leq \epsilon$ is called an ϵ-Nash equilibrium.

3 MODELING FOR HAND ABSTRACTION

In this section, we present a more precise modeling of the hand abstraction task. To this end, we first impose a series of constraints on IIGs to construct a subset of IIGs, namely signal observation ordered games (SOOGs), which is used to model hold'em games. We then model the hand abstraction task using signal observation abstraction within the SOOG framework.

3.1 Signal Observation Ordered games

Modeling hold'em games via traditional imperfect-information games has a critical limitation: infosets are treated as indivisible atomic units. In reality, a player's decision state in hold'em games inherently encompasses two independent core dimensions—chance action sequences (hand dealings) and rational players' action sequences. These are independent in that subsets can be defined for each separately, with any combination of one sequence from each subset forming a valid game history. Yet traditional IIGs merge these distinct components into a single, monolithic infoset, precluding independent analysis of each dimension, and leaving hand abstraction without a theoretical foundation. To address this gap, we extend the traditional IIG framework with additional constraints.

Before introducing SOOGs, we first define history-related operators on IIGs. For any $i \in \mathcal{N}_c$:

- **Trace extraction operator** $\mathcal{H}_{-i}(\cdot)$: Retains actions of all players except i in a history, replacing i's actions with wildcard $\emptyset_i$—a placeholder representing any possible action by player i. Example: For $h = a_i^1 \cdot a_j^1 \cdot a_k^1 \cdot a_i^2$ (where a_i^1, a_i^2 are i's actions, a_j^1, a_k^1 are actions of $j, k \in \mathcal{N}_c$) [2],
$$\mathcal{H}_{-i}(h) = \emptyset_i \cdot a_j^1 \cdot a_k^1 \cdot \emptyset_i.$$
- **Sequence extraction operator** $\tilde{\mathcal{H}}_{-i}(\cdot)$: Defined as $\tilde{\mathcal{H}}_{-i}(\cdot) := \mathcal{E} \circ \mathcal{H}_{-i}$, where $\mathcal{E}$ (wildcard elimination operator) removes all wildcards. For the above h:
$$\tilde{\mathcal{H}}_{-i}(h) = \mathcal{E}\left(\mathcal{H}_{-i}(h)\right) = \mathcal{E}\left(\emptyset_i \cdot a_j^1 \cdot a_k^1 \cdot \emptyset_i\right) = a_j^1 \cdot a_k^1.$$

We also define operators focusing on i:

- $\mathcal{H}_i(\cdot)$ (trace extraction for i): $\mathcal{H}_i(\cdot) := \bigcirc_{\substack{j \in \mathcal{N}_c \\ j \neq i}} \mathcal{H}_{-j}$, retaining i's actions with others replaced by wildcards [3].
- $\tilde{\mathcal{H}}_i(\cdot)$ (sequence extraction for i): $\tilde{\mathcal{H}}_i(\cdot) := \mathcal{E} \circ \mathcal{H}_i$, eliminating wildcards from $\mathcal{H}_i(h)$.

Example for h above:

$$\mathcal{H}_i(h) = a_i^1 \cdot \emptyset_j \cdot \emptyset_k \cdot a_i^2, \quad \tilde{\mathcal{H}}_i(h) = a_i^1 \cdot a_i^2.$$

---

[2] Omits initial h^o; full form: $h = h^o \cdot a_i^1 \cdot a_j^1 \cdot a_k^1 \cdot a_i^2$

[3] $\bigcirc$: sequential composition, $\bigcirc_{k=1}^N f_k := f_N \circ \cdots \circ f_1$

Definition 3.1 (Signal Observation Ordered Games (SOOG)). A signal observation ordered game is an imperfect-information game $\mathcal{G}$ (satisfying Definition 2.1) with the following additional constraints:

- $\gamma : H \mapsto \mathbb{N}^+$: A phase partition function. For any $h \in H$, $\gamma(h)$ denotes h's phase, counting chance histories ($h' \in H_c$) along the path from h^o to h (including h). Let $\Gamma = \max_{h \in H} \gamma(h)$ denote the final phase; notably, $\gamma(h^o) = 1$ (so h^o is a chance history).
- Chance actions reveal signals: Let $\Theta = \{\tilde{\mathcal{H}}_c(h) \mid h \in H\}$ be the total set of signals. Chance actions depend only on revealed signals: for any $h, h' \in H_c$, if $\tilde{\mathcal{H}}_c(h) = \tilde{\mathcal{H}}_c(h')$, then $A(h) = A(h')$ (identical legal actions) and $\zeta(h, a) = \zeta(h', a)$ for all $a \in A(h)$ (identical action probability distributions). Since the chance function ζ at chance histories only correlates with already dealt signals, and all legal actions a at chance histories serve to reveal new signals, we can construct a more focused function $\xi(\theta, \theta')$ to simplify and replace $\zeta(h, a)$—where $\theta \in \Theta$ is the revealed signal at the current chance history, and $\theta' \in \Theta$ is the new signal revealed by action a.
- Signal-action separability: For any histories $h_1, h_2, h_1' \in H$, if $\tilde{\mathcal{H}}_c(h_1) = \tilde{\mathcal{H}}_c(h_1')$ (i.e., they share identical chance action sequences) and $\mathcal{H}_{-c}(h_1) = \mathcal{H}_{-c}(h_2)$ (i.e., they share identical non-chance action traces), then there must exist a history $h_2' \in H$ such that $\tilde{\mathcal{H}}_c(h_2') = \tilde{\mathcal{H}}_c(h_2)$ and $\mathcal{H}_{-c}(h_2') = \mathcal{H}_{-c}(h_1')$.
- Signal observation partitions: Let $\Psi = (\Psi_i)_{i \in \mathcal{N}}$, where Ψ_i (signal observation infosets for i) is a partition of Θ. Let $\vartheta = (\vartheta_i)_{i \in \mathcal{N}}$ be a tuple of observation functions, where $\vartheta_i : \Theta \mapsto \Psi_i$ maps each signal $\theta \in \Theta$ to its corresponding signal observation infoset $\psi \in \Psi_i$. Signals in the same $\psi \in \Psi_i$ are indistinguishable to i.
- Phase-specific subsets: For each phase $r \in \{1, \ldots, \Gamma\}$, define: $H^{(r)} = \{h \in H \mid \gamma(h) = r\}$ (phase-r histories), $H_i^{(r)} = H^{(r)} \cap H_i$ (phase-r histories where i acts), $Z^{(r)} = Z \cap H^{(r)}$ (phase-r terminal histories), $\Theta^{(r)} = \{\tilde{\mathcal{H}}_c(h) \mid h \in H^{(r)}\}$ (phase-r signals), $\Psi_i^{(r)} = \{\psi \cap \Theta^{(r)} \mid \psi \in \Psi_i\}$ (phase-r signal observation infosets for i), $\Psi^{(r)} = (\Psi_i^{(r)})_{i \in \mathcal{N}}$ (phase-r signal observation partitions), $\mathcal{I}_i^{(r)} = \{I \in \mathcal{I}_i \mid I \subseteq H^{(r)}\}$ (phase-r infosets for i), and $\mathcal{I}^{(r)} = (\mathcal{I}_i^{(r)})_{i \in \mathcal{N}}$ (phase-r information partitions).
- History indistinguishability criterion: For any $h, h' \in H_i$ and $i \in \mathcal{N}$, h and h' belong to the same infoset $I \in \mathcal{I}_i$ if and only if $\mathcal{H}_{-c}(h) = \mathcal{H}_{-c}(h')$ (identical non-chance action traces) and $\vartheta_i(\tilde{\mathcal{H}}_c(h)) = \vartheta_i(\tilde{\mathcal{H}}_c(h'))$ (indistinguishable signals).
- Survival functions: $\omega = (\omega_i)_{i \in \mathcal{N}}$ is a tuple of survival functions, where for each $i \in \mathcal{N}$ and $h \in H$,

  $$\omega_i(h) = \mathbb{I}\{\text{player } i \text{ still participates at } h\},$$

  with the property that if $h \sqsubseteq h'$ and $\omega_i(h) = 0$, then $\omega_i(h') = 0$ (i.e., the exit from the game is irreversible). Here, $\mathbb{I}\{\cdot\}$ denotes the indicator function (1 if the condition holds, 0 otherwise).
- Terminal order and utility consistency: In the final phase Γ, each signal $\theta \in \Theta^{(\Gamma)}$ (final signals) induces a total order $\preceq_\theta$ over $\mathcal{N}$, which satisfies reflexivity ($i \preceq_\theta i$ for all i), totality (for any $i, j \in \mathcal{N}$, either $i \preceq_\theta j$ or $j \preceq_\theta i$), and transitivity (if $i \preceq_\theta j$ and $j \preceq_\theta k$, then $i \preceq_\theta k$). For any terminal history $z \in Z^{(\Gamma)}$ with $\theta = \tilde{\mathcal{H}}_c(z)$, if $\omega_i(z)\omega_j(z) = 1$ (both i and j survive at z) and $i \preceq_\theta j$, then $u_i(z) \leq u_j(z)$.

Compared to IIGs, SOOGs better capture the core features of hold'em games, with two key refinements aligned to the structure of hold'em games:

First, SOOG's **signal-action separability** matches "hands vs. betting actions" split in hold'em games. Take a HULH example: Suppose a history h at phase 2 (Flop in HULH) where player 1 holds $\mathsf{9}_d\mathsf{9}_c$, player 2 holds $\mathsf{J}_h\mathsf{Q}_h$, and the community cards are $\mathsf{A}_s\mathsf{2}_h\mathsf{3}_h$; we denote this signal as $\theta = \tilde{\mathcal{H}}_c(h) = (\mathsf{9}_d^1\mathsf{9}_c^1, \mathsf{J}_h^2\mathsf{Q}_h^2) \cdot (\mathsf{A}_s\mathsf{2}_h\mathsf{3}_h)$ [4]. SOOG's observation functions map this signal to players' private views: $\vartheta_1(\theta) = (\mathsf{A}_h^1\mathsf{K}_h^1) \cdot (\mathsf{A}_s\mathsf{2}_h\mathsf{3}_h)$ and $\vartheta_2(\theta) = (\mathsf{J}_h^2\mathsf{Q}_h^2) \cdot (\mathsf{A}_s\mathsf{2}_h\mathsf{3}_h)$. Crucially, the betting actions in history h (e.g., player 1 raises, player 2 calls in Preflop phase) form the non-chance trace $\mathcal{H}_{-c}(h)$, and separability ensures this trace can combine with other signals (e.g., replacing player 1's hand with $\mathsf{A}_h\mathsf{K}_h$ to form $\theta' = (\mathsf{A}_h^1\mathsf{K}_h^1, \mathsf{J}_h^2\mathsf{Q}_h^2) \cdot (\mathsf{A}_s\mathsf{2}_h\mathsf{3}_h)$) to yield a valid hold'em history—unlike IIGs, which treat hands and actions as an indivisible unit.

Second, SOOG's **total order on final signals** models hold'em's showdown. At showdown, rewards depend on hand strength, which SOOG encodes as a total order $\preceq_\theta$ over the final signal $\theta \in \Theta^{(\Gamma)}$ (e.g., player 1's stronger hand gives $2 \preceq_\theta 1$). Combined with survival functions (both players survive to showdown, so $\omega_1(z) = \omega_2(z) = 1$), SOOG's utility constraint $u_i(z) \leq u_j(z)$ for $i \preceq_\theta j$ directly reflects the effect of showdown rewards.

These two refinements reveal SOOG signals can be studied independently and have inherent quality hierarchies (even in non-final phases), enabling standalone research and classification as an abstraction technique for this class of imperfect-information games.

3.2 Signal Observation Abstraction

We can now model the hand abstraction task using the SOOG framework:

Definition 3.2. In a SOOG $\mathcal{G}$, $\alpha = (\alpha_1, \ldots, \alpha_N)$ is a signal observation abstraction profile, and Ψ_i^α denotes the set of abstracted signal observation infosets for $i \in \mathcal{N}$ under α. Each $\alpha_i : \Theta \mapsto \Psi_i^\alpha$ maps a signal $\theta \in \Theta$ to an abstracted signal observation infoset $\psi^\alpha \in \Psi_i^\alpha$. Furthermore, each ψ^α can be partitioned into finer signal observation infosets within Ψ_i.

Given a signal observation abstraction profile $\alpha = (\alpha_1, \ldots, \alpha_N)$ for a SOOG $\mathcal{G}$, a **(signal observation) abstracted game** $\mathcal{G}^\alpha$ is derived by substituting ϑ_i with α_i in $\mathcal{G}$; signal observation abstraction does not directly alter the game itself but modifies the structure of players' strategy spaces, and since infosets in SOOGs are defined exclusively based on players' signal observations (i.e., the original infoset partition $\mathcal{I}_i$ is induced by ϑ_i), replacing ϑ_i with α_i reshapes the information structure: the abstracted infoset partition $\mathcal{I}_i^\alpha$ for player i is defined such that two original infosets $I, I' \in \mathcal{I}_i$ belong to the same abstracted infoset $I^\alpha \in \mathcal{I}_i^\alpha$ if and only if there exist histories $h \in I$ and $h' \in I'$ satisfying $\alpha_i(\tilde{\mathcal{H}}_c(h)) = \alpha_i(\tilde{\mathcal{H}}_c(h'))$ and $\mathcal{H}_{-c}(h) = \mathcal{H}_{-c}(h')$, and notably, $\mathcal{I}_i^\alpha$ forms a partition of $\mathcal{I}_i$—every original infoset $I \in \mathcal{I}_i$ is a subset of exactly one abstracted infoset $I^\alpha \in \mathcal{I}_i^\alpha$ (i.e., $I \subset I^\alpha$); in the abstracted game $\mathcal{G}^\alpha$, each player $i \in \mathcal{N}$ selects an **abstracted strategy** $\sigma_i^\alpha \in \Sigma_i^\alpha$ (where Σ_i^α denotes the set of all such abstracted strategies for i), formally,

---

[4] h: hearts, d: diamonds, s: spades, c: clubs; superscripts ^{p} denote private cards (e.g., $\mathsf{9}_d^p\mathsf{9}_c^p$ = player p's 9s in diamonds/clubs).

$\sigma_i^\alpha = \{\sigma_i^\alpha(I^\alpha, \cdot) \mid I^\alpha \in \mathcal{I}_i^\alpha\}$ with $\sigma_i^\alpha(I^\alpha, \cdot) : A(I^\alpha) \to [0, 1]$ being a probability distribution over the legal actions at I^α (with $A(I^\alpha) = A(I)$ for all $I \subset I^\alpha$, as legal actions are invariant across original infosets merged into I^α) and satisfying $\sum_{a \in A(I^\alpha)} \sigma_i^\alpha(I^\alpha, a) = 1$ (normalization) for each $I^\alpha \in \mathcal{I}_i^\alpha$, and when all players select their abstracted strategies, they form an abstracted strategy profile $\sigma^\alpha = (\sigma_1^\alpha, \dots, \sigma_N^\alpha)$ with the set of all possible abstracted profiles denoted by $\Sigma^\alpha = \times_{i \in \mathcal{N}} \Sigma_i^\alpha$; crucially, any abstracted strategy $\sigma_i^\alpha \in \Sigma_i^\alpha$ can induce a corresponding original strategy $\sigma_i \in \Sigma_i$ in $\mathcal{G}$: for every original infoset $I \in \mathcal{I}_i$ with $I \subset I^\alpha$ (for some $I^\alpha \in \mathcal{I}_i^\alpha$), the action probability distribution of σ_i at the original infoset I inherits that of σ_i^α at its corresponding abstracted infoset I^α — i.e., $\sigma_i(I, a) = \sigma_i^\alpha(I^\alpha, a)$ for all $a \in A(I)$. A signal observation abstraction profile α is said to have **perfect-recall** if the abstracted game $\mathcal{G}^\alpha$ induced by α has perfect-recall; otherwise, it is called an **imperfect-recall** abstraction. It should be noted that arbitrary historical information discard represents merely an extreme form of imperfect-recall abstraction.

4 EVALUATION FOR HAND ABSTRACTION ALGORITHMS

With a suitable mathematical model in place, we can now construct evaluation methods for hand abstraction algorithms. To this end, this section proposes the **resolution bound** metric, establishes **potential-aware outcome isomorphism** to estimate the maximum performance limits of existing hand abstraction algorithms, and analyzes their inherent shortcomings.

4.1 Resolution Bound

A reasonable evaluation of signal observation abstractions involves identifying the strategy profile with the best performance in the strategy space they construct (e.g., in 2-player zero-sum scenarios, finding the abstraction with the lowest exploitability). Yet this performance-based indirect evaluation is impractical due to computational constraints. We thus need a framework starting directly from the abstractions themselves, with one intuitive direction being to assess their granularity.

Definition 4.1. In a SOOG, α_i and β_i are signal observation abstractions for player $i \in \mathcal{N}$. The refinement relationship between α_i and β_i is defined as follows: If, for any $\psi^\beta \in \Psi_i^\beta$, there exist one or more abstracted signal observation infosets in Ψ_i^α such that their union forms a partition of ψ^β, then α_i is said to refine β_i, denoted as $\alpha_i \sqsupseteq \beta_i$. Furthermore, α_i is a common refinement of multiple signal observation abstractions $\alpha_i^1, \dots, \alpha_i^m$ for player i if $\alpha_i \sqsupseteq \alpha_i^j$ for all $j \in \{1, \dots, m\}$.

It is hard to judge the quality of two signal observation abstractions without a refinement relationship: in some abstracted signal observation infosets, α_i is more refined than β_i, while in others, β_i is more refined than α_i. However, among those with a refinement relationship, the more refined one has inherent advantages. As Waugh et al. [24] proved this property for IIGs, we extend it to SOOGs (proof in Appendix B):

Theorem 4.2 (Adapted to SOOG from Theorem 3 in [24]). *Let α_i and β_i be signal observation abstractions for player $i \in \{1, 2\}$ in a 2-player zero-sum perfect-recall SOOG $\mathcal{G}$. Let $\alpha = (\alpha_i, \vartheta_{-i})$ and $\beta = (\beta_i, \vartheta_{-i})$ be perfect-recall signal observation abstraction profiles, with $\mathcal{G}^\alpha$ and $\mathcal{G}^\beta$ denoting their induced abstracted games respectively. If $\alpha_i \sqsupseteq \beta_i$, and σ^α and σ^β are the strategies mapped back to the original game from the Nash equilibria of $\mathcal{G}^\alpha$ and $\mathcal{G}^\beta$ respectively, then $\epsilon_i(\sigma^\alpha) < \epsilon_i(\sigma^\beta)$.*

Theorem 4.2 imposes several restrictive conditions: it requires player i's opponents adopt no abstraction, and the profiles α, β satisfy perfect-recall. However, in practice, we need not adhere strictly to all these constraints. Fundamentally, Theorem 4.2 gives key intuition: more granular abstractions tend to enable more competitive strategies when solving the abstracted game.

Since signal observation abstractions—i.e., hand abstraction in hold'em games—are all algorithmically generated, the concept of common refinement can be extended to evaluate these algorithms' performance. For a given SOOG $\mathcal{G}$, if a signal observation abstraction α_i exists such that, for any valid parameters (within the algorithm's designed parameter space), the signal observation abstraction α_i' generated by a signal observation abstraction algorithm is refinable by α_i (i.e., $\alpha_i \sqsupseteq \alpha_i'$), then α_i is the common refinement of the algorithm (rigorously, this is restricted to $\mathcal{G}$ and rational player i). The common refinement of a signal observation abstraction algorithm defines the upper bound of its ability to distinguish signal observation infosets, termed the **resolution bound**. From Theorem 4.2, we infer that high-quality solutions solved from the abstracted game induced by an observation abstraction algorithm's resolution bound tend to outperform those solved from the abstracted game induced by the algorithm's generated abstractions—though this performance advantage is not strictly guaranteed. Notably, a signal observation algorithm with a finer resolution bound does not guarantee better performance; however, identifying a coarse resolution bound for such an algorithm indicates inherent flaws in the algorithm.

4.2 Potential-Aware Outcome Isomorphism

We now construct a flawed signal observation abstraction for a given SOOG $\mathcal{G}$, and subsequently prove that this flawed abstraction serves as the resolution bound of existing popular signal observation abstraction algorithms that follow the paradigm of arbitrary historical information discard. Through this analysis, we aim to reveal the inherent defects of current algorithms and elaborate on the underlying causes of these limitations.

Without loss of generality, we construct the signal observation abstraction from the perspective of rational player i. We extend the nature's reaching probability notation $\pi_c(h)$ (defined for histories h) to signals θ: let $\pi_c(\theta) = \prod_{\theta' \sqsubseteq \theta} \xi(\theta'', \theta')$, where θ'' is θ''s immediate predecessor, and $\sqsubseteq$ denotes signal predecessor (or equal).

First, for any signal observation infoset $\psi \in \Psi_i^{(r)}$ in the final phase $r = \Gamma$, we define a winrate outcome feature as

$$wo_i^{(r)}(\psi) = \left(wo_i^{(r),0}(\psi), wo_i^{(r),1}(\psi), \dots, wo_i^{(r),N}(\psi)\right),$$

where we denote $\mathcal{N}_{-i} = \mathcal{N} \setminus \{i\}$. The components are defined as follows:

- $wo_i^{(r),0}(\psi)$ denotes the probability that player i ranks lower than at least one other player. Formally:

$$wo_i^{(r),0}(\psi) = \frac{\sum_{\theta \in \psi} \pi_c(\theta) \cdot \mathbb{I}\{\exists j \in \mathcal{N}_{-i} \text{ such that } j \not\preceq_\theta i \text{ and } i \preceq_\theta j\}}{\sum_{\theta \in \psi} \pi_c(\theta)}.$$

- For $l > 0$, $wo_i^{(r),l}(\psi)$ denotes the probability that player i ranks no lower than any other player and ranks higher than exactly $l - 1$ other players. Formally:

$$wo_i^{(r),l}(\psi) = \qquad (1)$$
$$\frac{\sum_{\theta \in \psi} \pi_c(\theta) \cdot \mathbb{I}\{\forall j \in \mathcal{N}_{-i}, j \preceq_\theta i \text{ and } |\{j \in \mathcal{N}_{-i} \mid i \preceq_\theta j\}| = l - 1\}}{\sum_{\theta \in \psi} \pi_c(\theta)}.$$

where $|\{\cdot\}|$ denotes the cardinality of the set (i.e., the count of elements satisfying the condition).

Specifically, in the 2-player scenario, only $wo_i^{(\Gamma),0}(\psi)$, $wo_i^{(\Gamma),1}(\psi)$, and $wo_i^{(\Gamma),2}(\psi)$ remain. In hold'em games, these correspond respectively to the losing rate, tying rate, and winning rate when enumerating the opponent's hands against the player's own hand.

Then, we construct a **potential-aware outcome feature (PAOF)** for each signal observation infoset in each phase. Signal observation infosets with identical PAOFs are grouped into the same class, while those with distinct PAOFs belong to different classes; the resulting classes are referred to as **potential-aware outcome isomorphisms (PAOIs)**. The number of distinct PAOIs in phase r is denoted as $C^{(r)}$.

First, for the final phase Γ, the PAOF is defined as:

$$pf^{(\Gamma)}(\psi) := wo_i^{(\Gamma)}(\psi).$$

Second, for a non-final phase r and a signal observation infoset $\psi \in \Psi_i^{(r)}$, the PAOF $pf^{(r)}(\psi)$ is a histogram of the PAOIs from the next phase. Specifically, let $pi_j^{(r+1)}$ denote the set of signal observation infosets contained in the j-th PAOI of phase $r + 1$. The j-th component of the PAOF of ψ is then given by:

$$pf_j^{(r)}(\psi) = \frac{\sum_{\theta \in \psi} \pi_c(\theta) \cdot \sum_{\psi' \in pi_j^{(r+1)}} \sum_{\theta' \in \psi'} \xi(\theta, \theta')}{\sum_{\theta \in \psi} \pi_c(\theta) \cdot \sum_{\theta' \in \Theta^{(r+1)}} \xi(\theta, \theta')}.$$

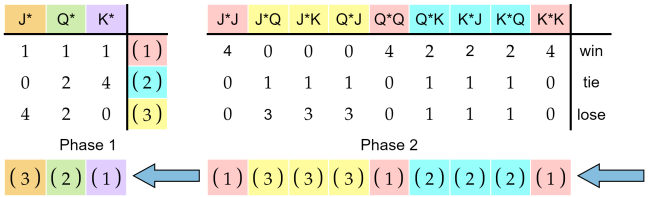

J*	Q*	K*	
1	1	1	(1)
0	2	4	(2)
4	2	0	(3)

J*J	J*Q	J*K	Q*J	Q*Q	Q*K	K*J	K*Q	K*K	
4	0	0	0	4	2	2	2	4	win
0	1	1	1	0	1	1	1	0	tie
0	3	3	3	0	1	1	1	0	lose

Figure 1: The process of constructing potential-aware outcome isomorphism in Leduc Hold'em.

Figure 1 illustrates PAOI construction in Leduc Hold'em (rules in Appendix A.2). Here, the first letter of each hand denotes a player's private card, and '*' represents the opponent's hidden card. In the second phase, the second letter indicates the revealed community card; each hand corresponds to 4 possible opponent card combinations, yielding 4 signals per infoset. After showdown, we group signal observation infosets into abstracted sets if they share identical win-tie-loss outcome distributions (labeled accordingly). Counts replace win rates here due to equal opponent combinations across hands.

In the first phase, each hand branches into second-phase abstracted signal observation infosets as 5 possible cards are dealt. We group hands with identical branching distributions into abstracted sets; counts again replace probabilities. Leveraging Leduc's perfect-recall (each infoset maps directly to next-phase hands), we enumerate infosets directly instead of individual signals. Notably, this bottom-up construction—where later-phase equivalence classes are built without reference to early-phase information—exemplifies the paradigm of arbitrary historical information discard.

Observations reveal that first-phase abstracted signal observation infosets match the original count (3, indicating strong discriminative power), while the second phase reduces to 3 abstracted sets—far fewer than the original 9—exposing PAOI's limited ability to distinguish infosets. Below are two propositions, with proofs in Appendix C.

Proposition 4.3. *For a hold'em game modeled as a SOOG $\mathcal{G}$, the PAOI abstraction serves as a resolution bound of algorithm EHS.*

Proposition 4.4. *For a hold'em game modeled as a SOOG $\mathcal{G}$, the PAOI abstraction serves as a resolution bound of algorithm PAAEMD.*

Propositions 4.3 and 4.4 suggest that the performance of strategies solved under EHS and PAAEMD abstractions will likely be weaker than that of their counterparts under PAOI (details of the EHS and PAAEMD algorithms are provided in the Appendix C) abstraction. Nevertheless, we next point out and analyze that PAOI itself is not an abstraction with strong performance.

Table 1 presents the number of distinct signal observation infosets across different phases in HULH (ground truth) — with more detailed rules available in Appendix A.1 — and compares it with those abstracted by LI and PAOI. Notably, PAOI's abstraction follow a spindle-shaped distribution: fewer in early/late phases but more in middle phases. By contrast, ground truth infosets exhibit a clear triangular pattern—gradually increasing with game progression—consistent with natural strategic information accumulation. The LI algorithm effectively preserves this pattern, aligning with the game's inherent information growth. PAOI's spindle-shaped distribution, however, deviates sharply from this baseline, implying systematic late-game information loss. This raises concerns about current outcome-based imperfect-recall hand abstraction algorithms (e.g., EHS and PAAEMD) using PAOI as their resolution bound, as such deviation could significantly compromise the effectiveness of their solved strategies—particularly in late phases.

To explain PAOI's systematic late-game information loss, we first identify its core limitation: it arbitrarily discards all historical information, classifying signal observation infosets solely by current and future phases. As shown in Figure 2, the number of abstracted signal observation infosets under PAOI is jointly determined by two key factors, whose interaction ultimately causes information loss. On one hand, inherent signal observation infosets increase with game progression, theoretically enabling more distinct classes. On the other hand, the effectiveness of PAOI's classification—i.e., distinguishing whether two signal observation infosets belong to the same abstract class—depends on the data volume used to compute their PAOFs. For example, in HULH, calculating the PAOF

Table 1: Quantity of signal observation infosets in unabstracted game and abstracted signal observation infosets for various algorithms in HULH.

Phase	No Abstraction	LI	PAOI
Preflop	$\binom{52}{2}$=1326	169	169
Flop	$\binom{52}{2,3}$=25989600*	1286792	1137132
Turn	$\binom{52}{2,3,1}$=1195521600	55190538	2337912
River	$\binom{52}{2,3,1,1}$=5379847200	2428287420	20687

* The notation $\binom{a}{b,c}$ denotes a combinatorial count: given a set of a elements in total, it first selects b elements without replacement, then selects c elements without replacement from the remaining $a - b$ elements. The total count is equivalent to $\binom{a}{b} \cdot \binom{a-b}{c}$.

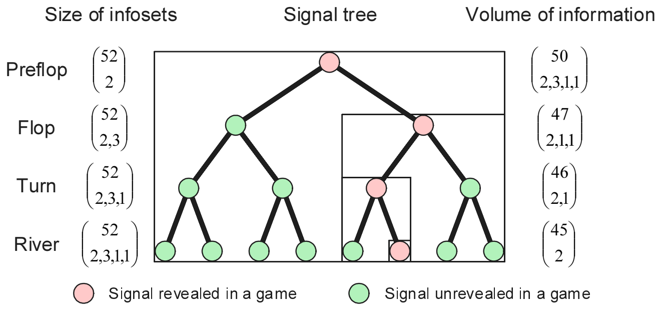

Figure 2: Two factors influencing the quantity of distinct abstracted signal observation infosets for algorithms that only use current and future information.

of a signal observation infoset (a hand) requires enumerating all possible rollout scenarios: in the River phase, this involves $\binom{52-7}{2}$ cases (opponent's hole cards), while in the Turn phase, it requires $\binom{52-6}{2,1}$ cases (opponent's hole cards and final River community card). As Figure 2 illustrates, the data volume needed to assign a signal (each pink node) to an abstract class equals the data volume of the subtree rooted at that node—and this volume declines progressively in later phases. Less data increases the likelihood of identical PAOFs for distinct signal observation infosets, erroneously grouping them into the same equivalence class. This not only reduces the number of distinct abstracted infosets far below the original count but also exacerbates **excessive abstraction** in late game, ultimately leading to the observed spindle-shaped distribution of PAOI classes.

4.3 Full-Recall Outcome Isomorphism

We have identified the underlying reasons for PAOI's excessive abstraction phenomenon. To address this, we improve PAOI by incorporating historical information and construct a new abstraction: **full-recall outcome isomorphism (FROI)**, a special case of **k-recall outcome isomorphism (k-ROI)**.

k-ROI is developed from PAOI. For any $\psi \in \Psi_i^{(r)}$ of phase r in a SOOG $\mathcal{G}$ (for player i), we construct its k**-recall outcome feature (k-ROF)** — a k-dimensional vector — defined as:

$$rf_i^{(r,k)}(\psi) = \left(pi^{(r)}(\psi^{(r)}), pi^{(r-1)}(\psi^{(r-1)}), \ldots, pi^{(r-k)}(\psi^{(r-k)})\right),$$

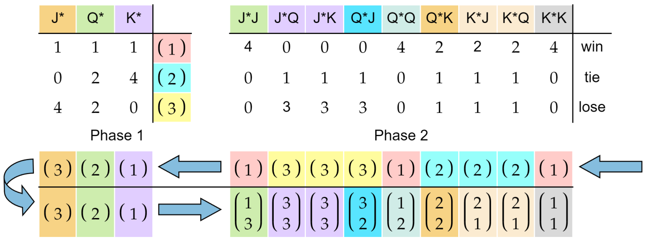

J*	Q*	K*	
1	1	1	(1)
0	2	4	(2)
4	2	0	(3)

J*J	J*Q	J*K	Q*J	Q*Q	Q*K	K*J	K*Q	K*K	
4	0	0	0	4	2	2	2	4	win
0	1	1	1	0	1	1	1	0	tie
0	3	3	3	0	1	1	1	0	lose

Figure 3: The process of constructing full-recall outcome isomorphism in Leduc Hold'em.

where $\psi^{(r')}$ (for $r' < r$) denotes the predecessor of ψ at phase r' (by the perfect-recall assumption, the predecessors of all signals in ψ at the same phase r' belong to the same signal observation infoset), and $pi^{(r')}(\psi^{(r')})$ represents the label (index) of the PAOI class to which $\psi^{(r')}$ belongs.

Figure 3 illustrates the construction of 0-ROI for phase 1 hands and 1-ROI for phase 2 hands in Leduc Hold'em. This is a two-step process: first, PAOFs and PAOIs for the current and all preceding phases are computed bottom-up; then, k-ROF and k-ROI for each hand are calculated top-down. For phase r, the $(r-1)$-ROI is referred to as **full-recall outcome isomorphism (FROI)**. Notably, "full-recall" here does not denote perfect-recall in the abstracted game; it only indicates inclusion of all prior phase information. This distinction is critical: future clustering-based abstraction algorithms targeting FROI as their resolution bound will likely remain imperfect-recall, yet they demonstrate a viable path to integrate historical context into imperfect-recall abstraction frameworks.

Table 2 reports the number of distinct abstracted classes identifiable by k-ROI (with $k = 0, \ldots, r-1$) in phase r of HULH. Clearly, as k increases (i.e., more historical information is incorporated), the number of distinct abstracted classes identifiable by k-ROI rises. Additionally, FROI's abstracted class counts across phases—169, 1241210, 42040233, and 638585633—exhibit the triangular pattern seen in both no abstraction and LI (Table 1).

5 EXPERIMENT

To provide direct evidence that abstraction methods arbitrarily discarding historical information suffer significant performance deficits, we evaluate the performance of ϵ-Nash equilibria solved in abstracted games when applied to the original game.

First, we select the experimental environment. Given HULH's extreme complexity, we adopt a simplified setting. Notably, conventional simplified environments (e.g., Kuhn Poker, Leduc Hold'em, Flop Hold'em) have no more than 2 phases. As Table 2 shows, even in HULH (a game with a vast range of hands), excessive abstraction is not prominent at the 2-phase stage (1137132 vs. 1241210). We therefore construct a 3-phase variant named Numeral211 Hold'em, with rules detailed in Appendix A.3. Table 3 reports the number of (abstracted) signal observation infosets identified by various methods (and the unabstracted setting) in Numeral211 Hold'em, revealing a substantial discrepancy between PAOI (100, 2250, 3957) and FROI (100, 2260, 51228)—particularly in phase 3.

We conduct two evaluations for each assessed set of signal observation abstractions, denoted as $\alpha = (\alpha_1, \alpha_2)$. First is **asymmetric abstraction**, following the setup in Theorem 4.2: we construct two

Table 2: The number of signal observation infoset equivalence classes identified by k-ROI in each phase and k of HULH.

Phase	Preflop 1	Flop 2		Turn 3			River 4			
Recall (k)	0	0	1	0	1	2	0	1	2	3
k-ROI	169	1137132	1241210	2337912	38938975	42040233	20687	39792212	586622784	638585633

Table 3: Quantity of signal observation infosets in unabstracted game and abstracted signal observation infosets for various algorithms in Numeral211 Hold'em.

	Phase 1	Phase 2		Phase 3		
No Abstraction	780	29640		1096680		
LI	100	2260		62020		
Recall	0	0	1	0	1	2
k-ROI	100	2250	2260	3957	51176	51228

abstractions $\alpha' = (\alpha_1, \vartheta_2)$ and $\alpha'' = (\vartheta_1, \alpha_2)$ (where ϑ_1/ϑ_2 denote no abstraction for player 1/2), use CFR to solve ϵ-Nash equilibria σ' and σ'' for abstracted games $\mathcal{G}^{\alpha'}$ and $\mathcal{G}^{\alpha''}$, then concatenate these into a joint strategy $\sigma = (\alpha'_1, \alpha''_2)$ and evaluate its exploitability in the original game $\mathcal{G}$. This method has theoretical guarantees—finer abstractions yield stronger-performing strategies—but it does not reflect typical hand abstraction scenarios: in large-scale games, solving strategies with no opponent abstraction is computationally infeasible due to massive space overhead. We therefore adopt a second scenario, **symmetric abstraction**: given α, we solve the ϵ-Nash equilibrium of the abstracted game $\mathcal{G}^{\alpha}$ and evaluate its exploitability in $\mathcal{G}$. This scenario is susceptible to **abstraction pathology** [24]—a rare phenomenon where finer abstractions may unexpectedly produce poorer-performing strategies.

In the first experiment, we verify that mainstream signal observation abstraction algorithms perform barely better than PAOI. Specifically, we use EHS and PAAEMD algorithms: we apply no abstraction in phase 1, and construct abstraction classes accounting for 1/10 of PAOI's maximum discriminative capacity in phases 2 and 3 (i.e., 100-225-396). Given clustering algorithms' sensitivity to initial values, we generate 5 independent groups of abstractions for each algorithm. Subfigures (a1) and (a2) of Figure 4 present results for asymmetric and symmetric abstraction scenarios, respectively; EHS and PAAEMD both perform significantly worse than PAOI—particularly in exploitability, where higher values indicate poorer performance.

In the second experiment, we compare performance across PAOI, FROI, and LI. Notably, LI is a lossless abstraction based on hand suit rotation and serves as ground truth—critically, it cannot be further compressed via clustering. Figure 4(b1, b2) shows results for asymmetric and symmetric abstraction: PAOI-derived solutions perform significantly worse than FROI and LI, while FROI's performance is nearly identical to LI's—especially in asymmetric abstraction. To distinguish this subtle gap, Figure 4(c1, c2) presents results on a logarithmic scale for clearer resolution of minor differences.

These results confirm that discarding historical information systematically degrades abstraction algorithm performance, whereas incorporating historical information raises the performance upper bound of such algorithms.

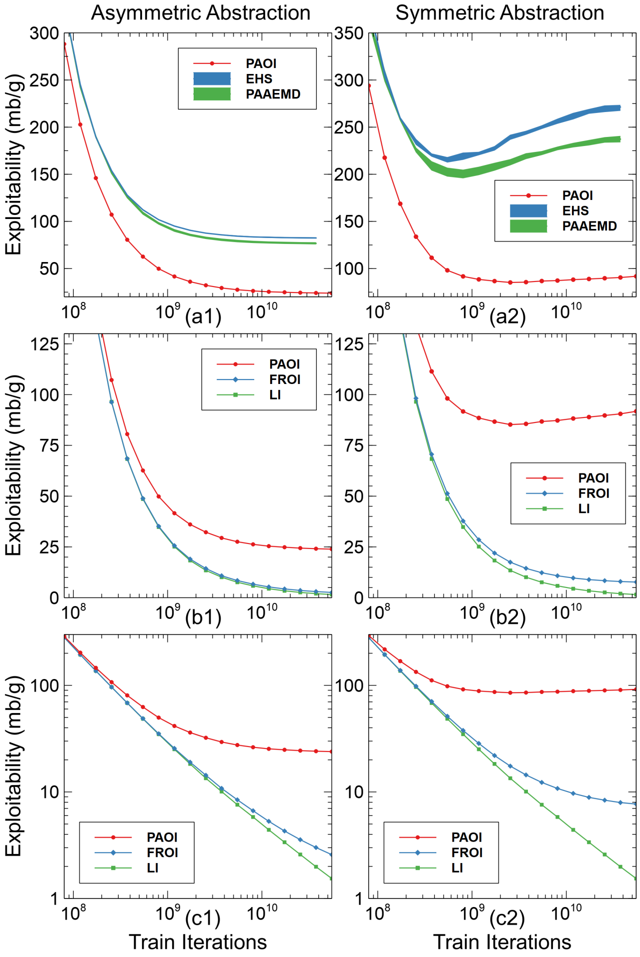

Figure 4: Experimental results on Numeral211 Hold'em.

6 CONCLUSION

We present a mathematical model and metric for hand abstraction, and show that PAOI acts as a low-resolution bound that degrades AI performance by discarding historical information. We then introduce FROI, which uses historical information to achieve a finer bound, with experiments validating its advantage. As FROI is an abstraction rather than an algorithm, future work will develop algorithms targeting the FROI bound.

ACKNOWLEDGMENTS

This work is jointly supported by the Key Research Project of Chinese Academy of Sciences (No.RCJJ-145-24-15)

REFERENCES

[1] D Billings, N Burch, A Davidson, R Holte, J Schaeffer, T Schauenberg, and D Szafron. 2003. Approximating game-theoretic optimal strategies for full-scale poker. In *International Joint Conference on Artificial Intelligence (IJCAI)*, Vol. 3. 661–668.

[2] Noam Brown and Tuomas Sandholm. 2017. Safe and nested subgame solving for imperfect-information games. *Advances in neural information processing systems* 30 (2017).

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

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

[5] Noam Brown and Tuomas Sandholm. 2019. Superhuman AI for multiplayer poker. *Science* 365, 6456 (2019), 885–890.

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

[7] Andrew Gilpin and Tuomas Sandholm. 2006. A competitive Texas Hold'em poker player via automated abstraction and real-time equilibrium computation. In *National Conference on Artificial Intelligence (NCAI)*, Vol. 21. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 1007.

[8] Andrew Gilpin and Tuomas Sandholm. 2007. Better automated abstraction techniques for imperfect information games, with application to Texas Hold'em poker. In *International Joint Conference on Artificial Intelligence (IJCAI)*. 1–8.

[9] Andrew Gilpin and Tuomas Sandholm. 2007. Lossless abstraction of imperfect information games. *Journal of the ACM (JACM)* 54, 5 (2007), 25–es.

[10] Andrew Gilpin, Tuomas Sandholm, and Troels Bjerre Sørensen. 2007. Potential-aware automated abstraction of sequential games, and holistic equilibrium analysis of Texas Hold'em poker. In *National Conference on Artificial Intelligence (NCAI)*, Vol. 22. Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, 50.

[11] Michael Johanson, Kevin Waugh, Michael Bowling, and Martin Zinkevich. 2011. Accelerating best response calculation in large extensive games. In *IJCAI*, Vol. 11. 258–265.

[12] Vojtěch Kovařík, Martin Schmid, Neil Burch, Michael Bowling, and Viliam Lisý. 2022. Rethinking formal models of partially observable multiagent decision making. *Artificial Intelligence* 303 (2022), 103645.

[13] Christian Kroer and Tuomas Sandholm. 2016. Imperfect-recall abstractions with bounds in games. In *ACM Conference on Economics and Computation (EC)*. 459–476.

[14] Christian Kroer and Tuomas Sandholm. 2018. A unified framework for extensive-form game abstraction with bounds. *International Conference on Neural Information Processing Systems (NeurIPS)* 31 (2018).

[15] Harold W Kuhn. 1950. A simplified two-person poker. *Contributions to the Theory of Games* 1 (1950), 97–103.

[16] Marc Lanctot, Kevin Waugh, Martin Zinkevich, and Michael Bowling. 2009. Monte Carlo sampling for regret minimization in extensive games. *Advances in neural information processing systems* 22 (2009).

[17] Viliam Lisý and Michael Bowling. 2017. Eqilibrium Approximation Quality of Current No-Limit Poker Bots.. In *AAAI Workshops*.

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

[19] Martin Schmid, Matej Moravčík, Neil Burch, Rudolf Kadlec, Josh Davidson, Kevin Waugh, Nolan Bard, Finbarr Timbers, Marc Lanctot, G Zacharias Holland, et al. 2023. Student of games: A unified learning algorithm for both perfect and imperfect information games. *Science Advances* 9, 46 (2023), eadg3256.

[20] Jiefu Shi and Michael L Littman. 2001. Abstraction methods for game theoretic poker. In *Computers and Games: Second International Conference, CG 2000 Hamamatsu, Japan, October 26–28, 2000 Revised Papers 2*. Springer, 333–345.

[21] Finnegan Southey, Michael Bowling, Bryce Larson, Carmelo Piccione, Neil Burch, Darse Billings, and Chris Rayner. 2005. Bayes' bluff: opponent modelling in poker. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence*. 550–558.

[22] Michal Šustr, Vojtěch Kovařík, and Viliam Lisý. 2019. Monte Carlo Continual Resolving for Online Strategy Computation in Imperfect Information Games. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*. 224–232.

[23] Oskari Tammelin. 2014. Solving large imperfect information games using CFR+. *arXiv preprint arXiv:1407.5042* (2014).

[24] Kevin Waugh, David Schnizlein, Michael Bowling, and Duane Szafron. 2009. Abstraction pathologies in extensive games. In *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, Vol. 2. 781–788.

[25] Kevin Waugh, Martin Zinkevich, Michael Johanson, Morgan Kan, David Schnizlein, and Michael Bowling. 2009. A practical use of imperfect recall. In *Symposium on Abstraction, Reformulation and Approximation (SARA)*.

[26] Yichi Zhou, Tongzheng Ren, Jialian Li, Dong Yan, and Jun Zhu. 2020. Lazy-CFR: fast and near-optimal regret minimization for extensive games with imperfect information. In *8th International Conference on Learning Representations (ICLR 2020)(virtual)*. International Conference on Learning Representations.

[27] Martin Zinkevich, Michael Johanson, Michael Bowling, and Carmelo Piccione. 2007. Regret minimization in games with incomplete information. *Advances in neural information processing systems* 20 (2007).