CSCE 631 — Algorithmic Game Theory

Lecture 15: Game Abstraction I

Alan Kuhnle

Summer 2026

Texas A&M University

Today's Plan

Motivation: Why Abstraction?

The Scale Problem in Real Games

- Heads-up no-limit Texas hold'em has $\sim 10^{161}$ game states
- Even heads-up *limit* hold'em: $\sim 3.16 \times 10^{17}$ states (Bowling et al., 2015)
- Compare: chess $\sim 10^{47}$, Go $\sim 10^{170}$
- No algorithm can enumerate or store strategies for the full game

Core Idea

Solve a *smaller*, approximate version of the game, then map the solution back to the original. This is **game abstraction**.

Abstraction as Lossy Compression

- An abstraction maps a large game G to a smaller game G'
- Analogous to lossy compression:
  - Original game G: full-resolution signal
  - Abstract game G': compressed representation
  - Strategy mapping: decoder
- The quality question: how much *strategic* information is lost?
- Unlike data compression, the loss metric is **exploitability**, not reconstruction error

Two Dimensions of Abstraction

Dimension	What is merged?	Effect
Information abstraction	Information sets	Fewer decision points
Action abstraction	Available actions	Smaller action space

- Most practical systems use *both* simultaneously
- Information abstraction: "these two hands are strategically similar"
- Action abstraction: "only consider bets of $0.5\times$, $1\times$, $2\times$ pot"

Formal Framework

Extensive-Form Games: Recap

An extensive-form game Γ consists of:

- Players $N = \{1, \ldots, n\}$ plus chance player c
- Game tree with nodes $\mathcal{H}$ (histories), terminal nodes $\mathcal{Z}$
- Information sets $\mathcal{I}_i$ for each player i
- Action sets $A(h)$ at each node h
- Utility functions $u_i : \mathcal{Z} \to \mathbb{R}$

A **behavioral strategy** σ_i assigns a probability distribution over $A(I)$ at each information set $I \in \mathcal{I}_i$.

Definition: Game Abstraction

Given an extensive-form game $\Gamma = (N, \mathcal{H}, \mathcal{Z}, \mathcal{I}, A, u)$, an **abstraction** produces a game $\Gamma' = (N, \mathcal{H}', \mathcal{Z}', \mathcal{I}', A', u')$ together with a mapping $\phi : \Gamma \to \Gamma'$ such that:

1. $|\mathcal{I}'| \leq |\mathcal{I}|$ (information abstraction)
2. $|A'(I')| \leq \max_{I \in \phi^{-1}(I')} |A(I)|$ (action abstraction)

The mapping ϕ merges information sets and/or removes actions. A strategy σ' in Γ' induces a strategy σ in Γ via the inverse mapping ϕ^{-1}.

Lossless vs. Lossy Abstraction

Lossless Abstraction

An abstraction is **lossless** if a Nash equilibrium in Γ' maps to a Nash equilibrium in Γ with the same exploitability.

Lossy Abstraction

A **lossy** abstraction may introduce additional exploitability: a Nash equilibrium in Γ' may be exploitable in Γ.

- Lossless abstractions exist when merged information sets are *strategically equivalent* (same optimal play)
- In practice, nearly all useful abstractions are lossy
- The art: minimize exploitability while maximizing compression

Exploitability as Quality Metric

Definition: Exploitability

The **exploitability** of a strategy σ_i in a two-player zero-sum game is:

$$\text{expl}(\sigma_i) = \max_{\sigma_{-i}} u_{-i}(\sigma_i, \sigma_{-i}) - v^*_{-i}$$

where v^*_{-i} is the value of the game to player $-i$. Equivalently, it measures how much a best-response opponent can gain beyond the Nash value.

- A Nash equilibrium has exploitability $= 0$
- An abstraction-based strategy has exploitability ≥ 0 in the original game
- Goal: find abstractions that yield low exploitability in Γ

Information Abstraction

Information Abstraction: Merging Information Sets

- **Key idea**: group "similar" information sets into a single abstract information set
- Player must play the same strategy across all merged information sets
- In poker: group hands that have similar strategic properties

Example: Simplified Poker

Consider a game with cards $\{1, 2, \ldots, 10\}$.

- Full game: 10 information sets at the first decision point
- Abstraction: group into buckets $\{1, 2, 3\}$, $\{4, 5, 6, 7\}$, $\{8, 9, 10\}$
- Abstract game: 3 information sets at the first decision point

Signal-Based Abstraction

- In imperfect-information games, players receive **signals** (e.g., private cards)
- Two signals are strategically similar if they lead to similar *distributions over outcomes*
- Example: in Texas hold'em after the flop
  - $A\spadesuit K\spadesuit$ on a board with $Q\spadesuit J\spadesuit 3\heartsuit$
  - $A\heartsuit K\heartsuit$ on a board with $Q\heartsuit J\heartsuit 3\spadesuit$
  - These are *isomorphic* — identical strategic situation
- Beyond isomorphism: group signals with similar *equity distributions*

Hand Strength and Equity

Definitions

- **Hand strength (HS)**: probability of winning at showdown against a uniformly random opponent hand

$$\mathrm{HS}(h) = \Pr[\text{win} \mid \text{hand } h, \text{board } b]$$

- **Equity**: expected share of the pot, accounting for future community cards
- **Hand potential**: probability that a currently losing hand improves to win (or vice versa)

Grouping by hand strength alone is crude: it ignores *how* a hand might win or lose on future streets.

Equity Distributions

A richer representation: the **equity distribution** of a hand.

For hand h on board b with future cards to come:

$$F_{h,b}(x) = \Pr[\text{equity against random opponent} \leq x]$$

- This is a CDF over $[0, 1]$
- Two hands with the same equity distribution are strategically interchangeable (same distribution over outcomes against all opponent ranges)
- Grouping hands by equity distribution is more refined than grouping by a single hand-strength number

Earth Mover's Distance for Card Grouping

Earth Mover's Distance (EMD)

Given two distributions P and Q over a metric space (M, d):

$$\text{EMD}(P, Q) = \inf_{\gamma \in \Pi(P,Q)} \int_{M \times M} d(x, y)\, d\gamma(x, y)$$

where $\Pi(P, Q)$ is the set of all couplings of P and Q. Also known as the **Wasserstein-1 distance**.

- For discrete equity histograms: EMD reduces to optimal transport between histograms
- EMD respects the *ordering* of equity buckets (unlike KL divergence or total variation)
- Used by Johanson et al. (2013): cluster hands by EMD between their equity distributions

EMD-Based Clustering: Example

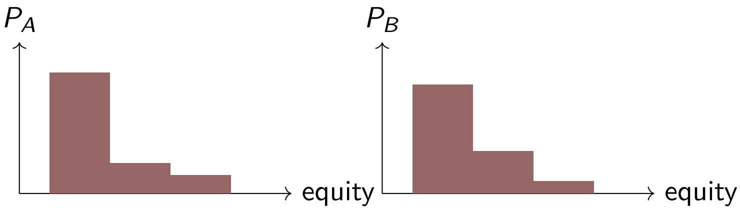

Hand A: weak draw

Hand B: similar draw

- EMD(P_A, P_B) is small $\Rightarrow$ group into same bucket
- Clustering algorithm: k-means with EMD as distance metric
- Typical: 169 preflop hand classes $\rightarrow$ 8–20 buckets

Action Abstraction

Action Abstraction: Reducing the Bet Space

- In no-limit poker, a player can bet any integer amount from 1 chip to their entire stack
- With 200 big blinds deep: ~ 200 possible bet sizes at each node
- **Action abstraction**: restrict to a small set of bet sizes

Common Bet Abstractions

Abstraction	Allowed bets
Coarse	fold, call, pot, all-in
Medium	fold, call, $\frac{1}{2}$pot, pot, $2\times$pot, all-in
Fine	fold, call, $\frac{1}{4}$pot, $\frac{1}{2}$pot, $\frac{3}{4}$pot, pot, $1.5\times$pot, $2\times$pot, all-in

The Granularity Trade-off

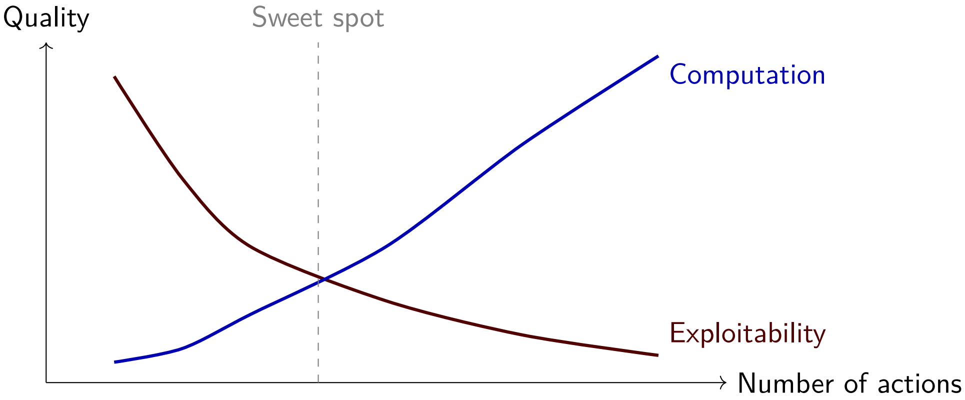

- More actions ⇒ lower exploitability but exponentially larger game tree
- Optimal granularity depends on available computation

Abstraction Pathologies

Abstraction Pathologies: A Cautionary Tale

Waugh et al. (2009)
Finer abstractions can produce *more exploitable* strategies than coarser ones.

This is counterintuitive:

- Intuitively, more detail should help
- But the equilibrium in a finer abstraction is optimized for the *abstract* game, not the original
- The abstraction boundary creates “seams” that an opponent in the real game can exploit

Pathology Example: Coin Game

Consider a simple game:

- Nature deals a coin: Heads (H) or Tails (T) with equal probability
- Player 1 sees the coin; Player 2 does not
- Player 1 chooses: Bet or Check
- If Bet: Player 2 chooses Call or Fold
- Payoffs depend on the coin

Coarse abstraction: merge H and T for Player 1 (single info set).
Fine abstraction: keep H and T separate.

Paradoxically, the coarse abstraction can yield a strategy that is *less exploitable* in certain game parameterizations, because the fine abstraction's equilibrium may create exploitable betting patterns.

Why Pathologies Occur

1. **Boundary effects**: the abstraction partitions hands into buckets, but hands near bucket boundaries are treated identically to hands far from the boundary
2. **Strategy mismatch**: the equilibrium in Γ' balances strategies across abstract information sets, but when mapped back to Γ, the balance is off
3. **Monotonicity failure**: the solution quality is *not* monotone in abstraction refinement

Implication

We cannot blindly refine an abstraction and expect improvement. Abstraction design requires careful analysis, not just increased resolution.

Theorem (Kroer & Sandholm, 2014)

For a two-player zero-sum extensive-form game, if the abstraction Γ' is obtained by merging information sets, the exploitability of the mapped strategy satisfies:

$$\mathrm{expl}_{\Gamma}(\sigma) \leq \mathrm{expl}_{\Gamma'}(\sigma') + \sum_{I \in \mathcal{I}} w_I \cdot d(I, \phi(I))$$

where $d(I, \phi(I))$ measures the "distance" between the original information set and its abstract image, and w_I are path weights.

- This bound is *loose* in practice, but provides theoretical guidance
- Minimizing the sum motivates EMD-based clustering

Strategy Mapping

From Abstract Strategy to Real Strategy

1. Solve Γ' to obtain equilibrium strategy σ'
2. Map σ' back to Γ using ϕ^{-1}:
   - For information abstraction: all original information sets mapped to the same abstract set play the same strategy
   - For action abstraction: actions not in the abstraction are unavailable — need **action translation**

Action Translation Problem

If the opponent plays an action not in the abstraction (e.g., bets 0.7$\times$ pot when only 0.5$\times$ and 1$\times$ are in the abstraction), how do we respond?

Action Translation Approaches

1. **Nearest action**: map to the closest action in the abstraction
   - Simple but introduces exploitable patterns
   - E.g., bet $0.7\times$ pot $\rightarrow$ round to $0.5\times$ or $1\times$
2. **Randomized mapping**: probabilistically map to neighboring abstract actions
   - Bet $0.7\times$ pot $\rightarrow$ play $0.5\times$ response with probability 0.6, play $1\times$ response with probability 0.4
   - Reduces exploitability from deterministic rounding
3. **Pseudo-harmonic mapping** (Schnizlein et al., 2009):
   - Weight abstract actions to preserve pot odds
   - Provably less exploitable than nearest action

Simplified Poker Examples

Example: Kuhn Poker

Kuhn Poker (Kuhn, 1950): the simplest nontrivial poker game.

- 3-card deck: $\{J, Q, K\}$
- Each player receives one card
- Betting round: check/bet, then call/fold
- 12 game states, 6 information sets per player

Abstraction exercise: merge J and Q for Player 1

- Reduces from 6 to 4 information sets
- The Nash equilibrium in the abstract game: Player 1 bets with K always, never with $\{J, Q\}$
- In the original game, optimal play bets with K always and bluffs with J at some frequency — this nuance is lost

Example: Leduc Hold'em

Leduc Hold'em (Southey et al., 2005): a toy poker game used as a standard benchmark.

- 6-card deck: $\{J, Q, K\}$ with 2 suits each
- Two rounds: one private card, then one public card
- Fixed bet sizes: 2 chips round 1, 4 chips round 2
- Maximum two raises per round
- ~ 936 information sets per player

Common abstractions studied:

- Merge suits (lossless — suits are symmetric)
- Merge J and Q in round 2 when board is K (lossy)
- The quality of these abstractions has been extensively characterized

Abstraction Quality in Leduc Hold'em

Abstraction	Info sets	Exploitability (mbb/h)
Full game (NE)	936	0
Suit-merged (lossless)	468	0
3-bucket preflop	312	42
2-bucket preflop	208	118
No card abstraction (single bucket)	104	350

- Exploitability measured in milli-big-blinds per hand (mbb/h)
- Suit-merging is lossless: an example of exploiting game symmetry
- Coarser card abstractions increase exploitability substantially

Summary

Summary

1. **Motivation**: real games are too large to solve; abstraction is necessary lossy compression
2. **Two types**: information abstraction (merge info sets) and action abstraction (reduce actions)
3. **Formal framework**: abstract game Γ', mapping ϕ, strategy translation ϕ^{-1}
4. **Quality metric**: exploitability of the mapped strategy in the original game
5. **Card grouping**: equity distributions + Earth Mover's Distance
6. **Pathologies**: finer abstraction does not guarantee better strategies (Waugh et al., 2009)
7. **Action translation**: mapping off-abstraction actions is a nontrivial problem

Next lecture: automated abstraction algorithms, safe subgame solving, and the connection to Libratus/Pluribus.