MIT 6.S890 — Team Leduc Poker Challenge

Instructor: Gabriele Farina

Fall 2024

CONTENTS Last update `2024-10-08`

1. Overview

Cooperation in gaming is always an important topic for discussion. It is known that when team members can communicate during play, the team can function as a single player with perfect recall, allowing the game to be solved using conventional methods for extensive-form games. However, when team members cannot communicate privately during play, the asymmetry in their observations prevents the application of standard approaches, presenting a challenging task.

The only permissible form of communication is team-private randomness before the game begins. Consequently, team members can agree on a correlated strategy that extracts external randomness. This correlated strategy can be viewed as a blend of multiple uncorrelated strategies, where team members act independently but their actions are coordinated based on a global signal. The primary objective of this project is to find an effective equilibrium in this class for poker, specifically focusing on Leduc Hold'em.

2. Game Rules

Team Leduc Hold'em is a simplified version of Limit Texas Hold'em, played with a reduced deck consisting only of six cards: Jacks (`J`), Queens (`Q`), and Kings (`K`) across two suits. The game progresses through two betting rounds. Initially, each player is dealt one private card. Following this, a single community card is revealed in the next betting round. Players are permitted one fixed bet per round. The game commences with each player placing a $1 ante into the pot. Players then take turns acting in sequence. All actions taken by any player are disclosed to all participants.

Players have three options during each round: call (or check), raise, and fold:

- **Call:** Players match the highest existing bet to remain in the hand. If no bets have been placed during the round, players may check, which allows them to stay in the game without increasing the bet.
- **Raise:** Players increase the pot by a pre-determined amount, specific to the betting round: The raise amount in the first betting round is $2, while the in the second betting round, the raise amount is $4. Each raise

adds to the total **bet** required to stay in the hand. Subsequent players must either call by matching this new total bet or raise further by the fixed amount.

- **Fold:** Players may opt out of the current hand if a raise is too high, forfeiting any claim to the pot and ending their participation in that round. This action prevents any further financial commitment in the current game.

The betting rounds end when all remaining players have matched the most recent bet, or all remaining players check. If at least two players remain after the second betting round, all must reveal their hands. The winner is determined by the highest-ranking hand, with ties split evenly among those tied. Hand strength is evaluated based on matching ranks between the private and community cards, then ordered from highest to lowest as King > Queen > Jack. The game's outcome is measured by the net monetary gains or losses for each player.

In this project, you will explore a 4-player variant of Team Leduc Hold'em. In each round, Player 1 acts first, followed by Player 2, and so on. Players 0 and 2 form one team, while Players 1 and 3 form the opposing team. The utility of each team is calculated as the sum of the utilities of its individual players. Communication among team members during gameplay is not allowed. Therefore, the information set for Team Leduc Hold'em remains the same as in the common Leduc Hold'em with multiple individual players. However, teams are allowed to use correlated strategies to improve their performance. Specifically, players can coordinate using a method of public randomness to plan and synchronize their strategies before each game session, aiming to maximize their profits.

3. Infoset Encoding

In Leduc Hold'em, the information set encodes all previous actions from all players, as well as the private card in the first betting round and the community card in the second betting round, if revealed. The sequence of each betting round is denoted by `/C:<C>/P<P1>:<A1>/P<P2>:<A2> .../`, where:

- `C` represents the card(s) revealed: the private card in the first round and the community card in the second round. Cards are represented by `J`, `Q`, or `K`, ignoring suits.
- Each `Pi Ai` pair represents an action by Player `Pi`. Here, `Ai` is a character that specifies the action performed by the player, where `c` stands for call or check, `r` for raise, and `f` for fold.

The information set of the player is given by concatenating the sequence from all previous betting rounds. For example, the string `/C:???K/P1:r/P2:f/P3:2f/P4:c/C:J/P1:c/` illustrates an information set from Player 4's perspective, who holds the private card `K`. The sequence unfolds with Player 1 raising, Player 2 folding, Player 3 calling, and Player 4 folding. Upon revealing the community card `J` in the second round, Player 1 checks. It is now Player 4's turn, following Player 2 and Player 3's earlier fold.

4. Input and Output

To ensure compability of any solution, we request that you submit strategy profiles computed by your algorithm. In the provided materials, you will find four files named `leduc4-infosets-player{i}.csv` for each i in $\{1, 2, 3, 4\}$. These files list all possible information sets for all four players.

Your submissions should include a meta strategy profile named `meta-strategy.csv` along with multiple strategy profiles named `strategy{j}-player{i}.csv` to encode your correlated strategy. The `meta-strategy.csv` file should contain several lines, each with two comma-separated values: j and p_j, which demonstrate the probability that signal j is selected with probability p_j.

Each strategy profile `strategy{j}-player{i}.csv` should encode the strategy of Player i when signal j is drawn. These strategy profiles must align with the line numbers of the corresponding infoset files. Each line in the strategy profiles should contain 3 comma-separated values, representing the probabilities of call (check) `c`, raise `r`, or fold `f`, respectively. Each value should be a floating-point number between 0 and 1, summing to 1 per line, and set to 0.0 for invalid moves.

We will provide sample policies, including a meta strategy profile named `example-meta-strategy.csv` and individual strategy profiles `example-strategy0-player{i}.csv`. These samples demonstrate a strategy where each player independently plays a valid move with equal probability in every information set.

5. Sample Gameplay

Here, we present a sample gameplay:

- At the start of the first betting round, each player contributes $1 as an ante. The cards dealt are `J`, `Q`, `Q`, and `K`, respectively.
  - Player 1, holding a `J`, initiates the betting with no prior actions, thus their information set is `/C:J???/`. They choose to raise, increasing the bets in the pot to $3, $1, $1, and $1 for each player respectively.
  - Player 2, with a `Q` and following Player 1′s raise, has an information set of `/C:?Q??/P1:r/`. They decide to fold.
  - Player 3, also with a Q, faces the action sequence `/C:??Q?/P1:r/P2:f/`. They too decide to fold.
  - Player 4, holding a K, has the information set `/C:???K/P1:r/P2:f/P3:f/`. They decide to call Player 1′s raise, adjusting the bets in the pot to $3, $1, $1, and $3 respectively.
- With all players either folding or matching the highest bet, the first betting round concludes. Only Players 0 and 3 remain. The community card dealt is a `J`.
  - Player 1, viewing no new actions, possesses the information set `/C:J???/P1:r/P2:f/P3:f/P4:c/C:J/`. They choose to check.
  - Player 4, facing the sequence `/C:???K/P1:r/P2:f/P3:f/P4:c/C:J/P1:c/` decides to raise. The raise amount in the second betting round is $4, changing the bets to $3, $1, $1, and $7 respectively.
  - Player 1, now facing `/C:J???/P1:r/P2:f/P3:f/P4:c/C:J/P1:c/P4:r/`, opts to re-raise to $11 (since Player 4′s last bet totaled $7), adjusting the stakes to $11, $1, $1, and $7 respectively.
  - Player 4, assessing the history `/C:???K/P1:r/P2:f/P3:f/P4:c/C:J/P1:c/P4:r/P1:r/`, decides to call Player 1′s bet, equalizing the bets in the pot at $11, $1, $1, and $11 respectively.
- As all active players have matched the latest bet, the second betting round ends. The two remaining players reveal their hands. Player 1′s private card `J` matches the community board `J`, securing them the win of the total pot of $24. Having contributed $11 to the pot, Player 1′s net gain is $13. Consequently, the utilities for the players are $13, −$1, −$1, and −$11, respectively. In terms of Team Leduc Hold'em, Player 1 and Player 3′s team achieves a net profit of $12, while Player 2 and Player 4′s team incurs a net loss of $12.

6. Note

There are about 148 thousand distinct infosets for each player in Team Leduc Hold'em. The solution concept related to the problem is TMECor. A possible algorithm for this problem is the team belief DAG [1], but several other possibilities exist.

Bibliography

[1] B. H. Zhang, G. Farina, and T. Sandholm, "Team belief DAG: generalizing the sequence form to team games for fast computation of correlated team max-min equilibria via regret minimization," in *International Conference on Machine Learning*, 2023, pp. 40996–41018.

6.S890 Topics in Multiagent Learning

Fall 2024