

IT3105 – Artificial Intelligence Programming

Kai Olav Ellefsen

Topics

- Summary from last time
- Opponent modelling
- Phases and deliverables

Project 1:

TEXAS HOLD'EM – SUMMARY SO FAR



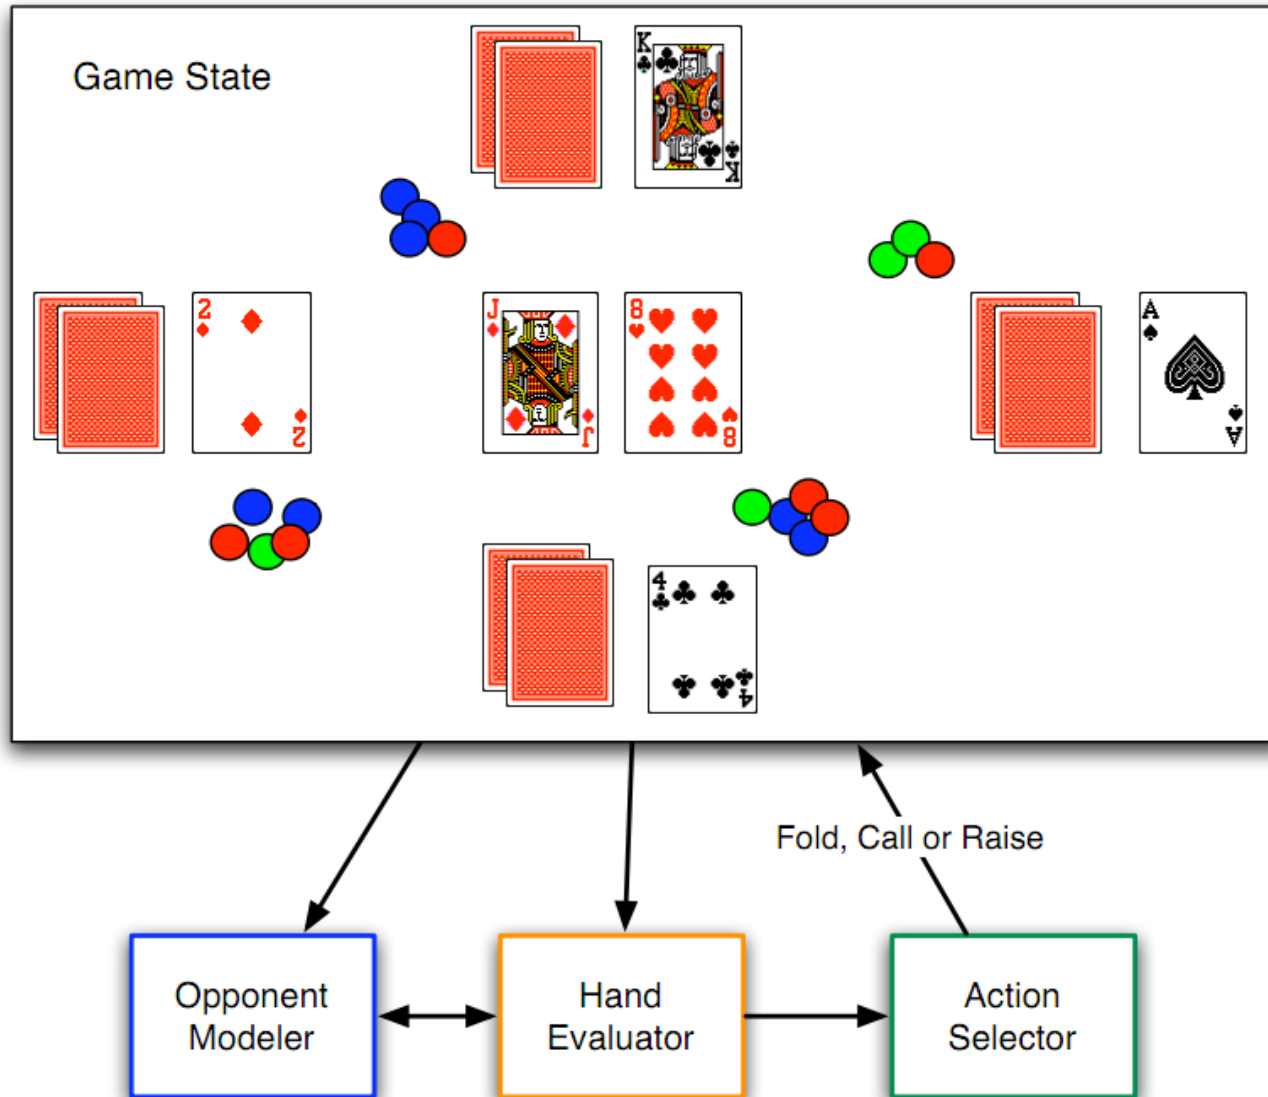
Source: Wikimedia Commons (By Todd Klassy)

URL: <http://commons.wikimedia.org/wiki/File:Holdem>

Game Classification


Imperfect Information	Risk Scotland Yard Battleship	Bridge Poker Hearts Scrabble
	Perfect Information	Chess Checkers Othello Go
	Deterministic	Stochastic

Elements of an AI Poker Player






Random Class Prototype



Wins
Ties
Losses

Income Rate
for class =
5&J (unsuited)
in a 4-player game

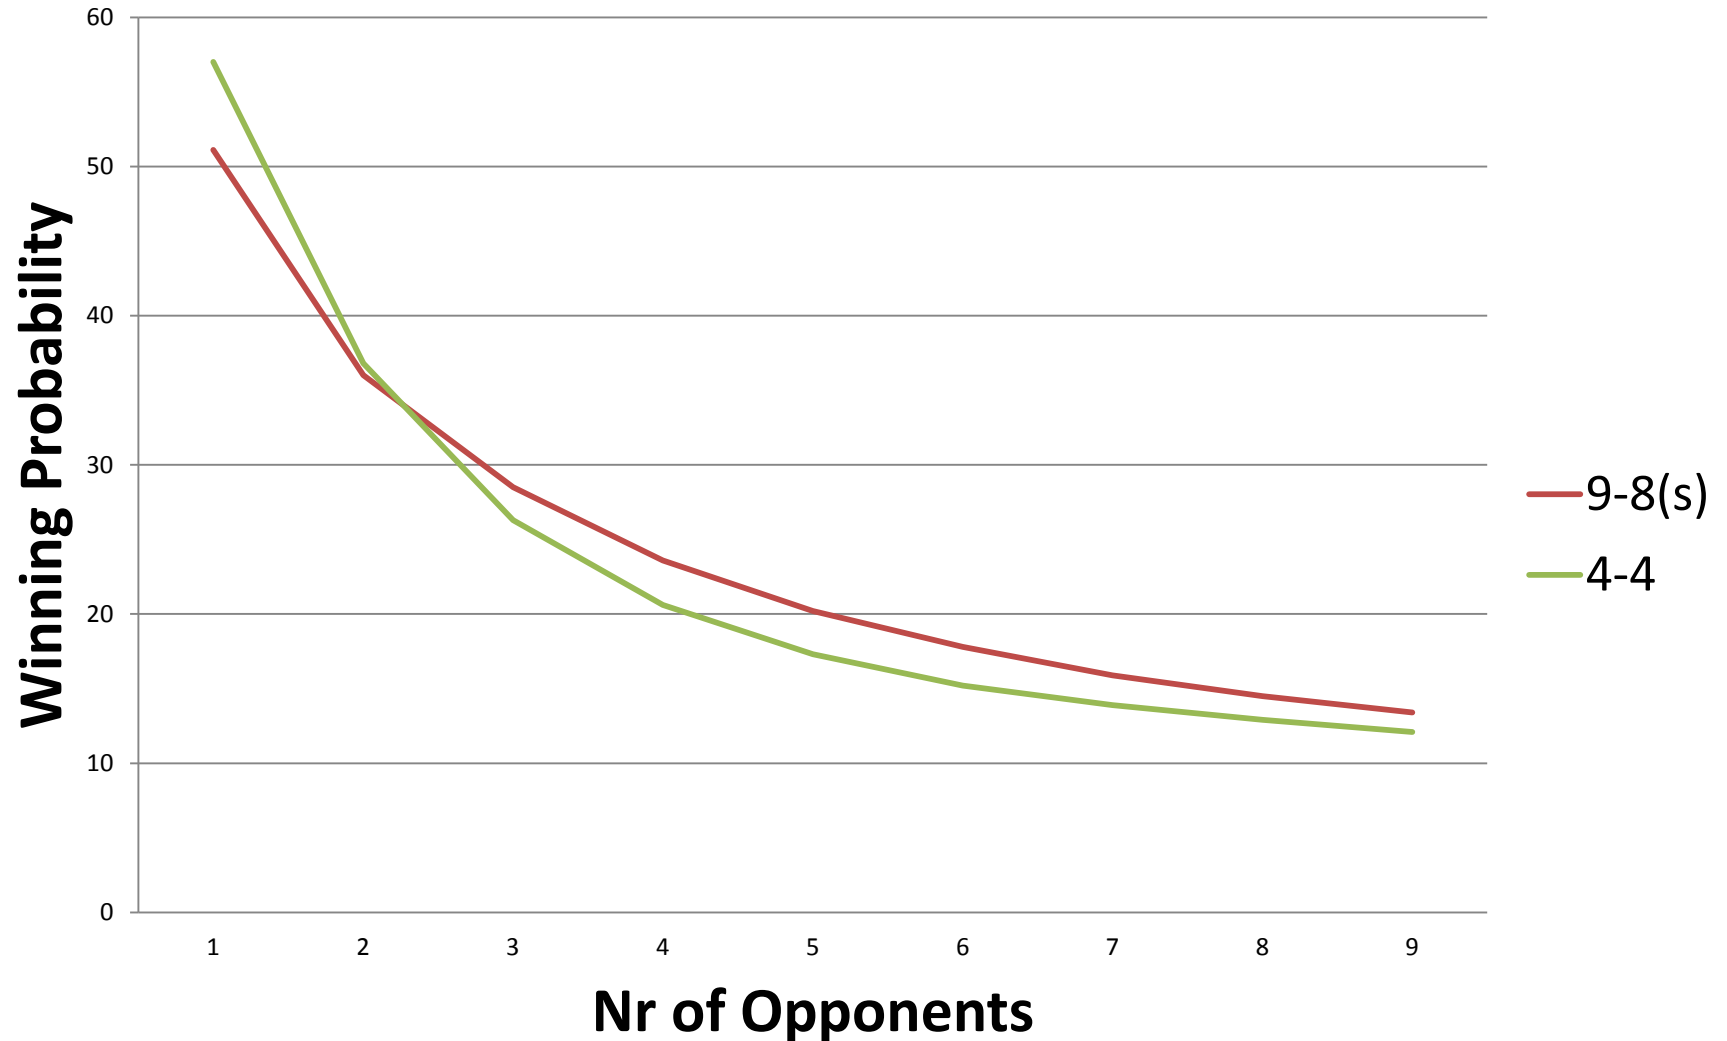
Equivalence Class Table



Interesting Question

- Do we really need to calculate data for different numbers of opponents?
 - A-A is always better than 2-7, right?
- Maybe we can just discount it?
 - $\text{WinProb} = (k/\text{nrOfPlayers}) * \text{defaultProb}$

9-8 Suited vs 4-4



Hand strength






Missing Turn and/or River cards are NOT filled in during the basic hand-strength calculation.



- vs -



$\binom{47}{2}$ possible pairs



hand strength = $\left(\frac{Wins + \frac{Ties}{2}}{Wins + Ties + Losses} \right)^k$



Hand Strength Example

- $$\text{hand strength} = \left(\frac{\text{Wins} + \frac{\text{Ties}}{2}}{\text{Wins} + \text{Ties} + \text{Losses}} \right)^k$$
- Against 1 opponent:
 - Hand strength is 0.585^1
- Against 5 opponents:
 - Hand strength is $0.585^5 = 0.069$
- Limitations:
 - Hand strength assumes *all players* reach showdown
 - Does not consider possible turn and river cards

Deciding on action

- Hand strength
- Pot odds
- Number of active opponents
- Number of raises
- Betting round (pre-flop, flop, turn, river)
- Opponent modelling

OPPONENT MODELLING



Source: Flickr (by [Jim Sher](#))

URL: <http://www.flickr.com/photos/blyzz/2545300865/in/phot>

Why Model Opponents?

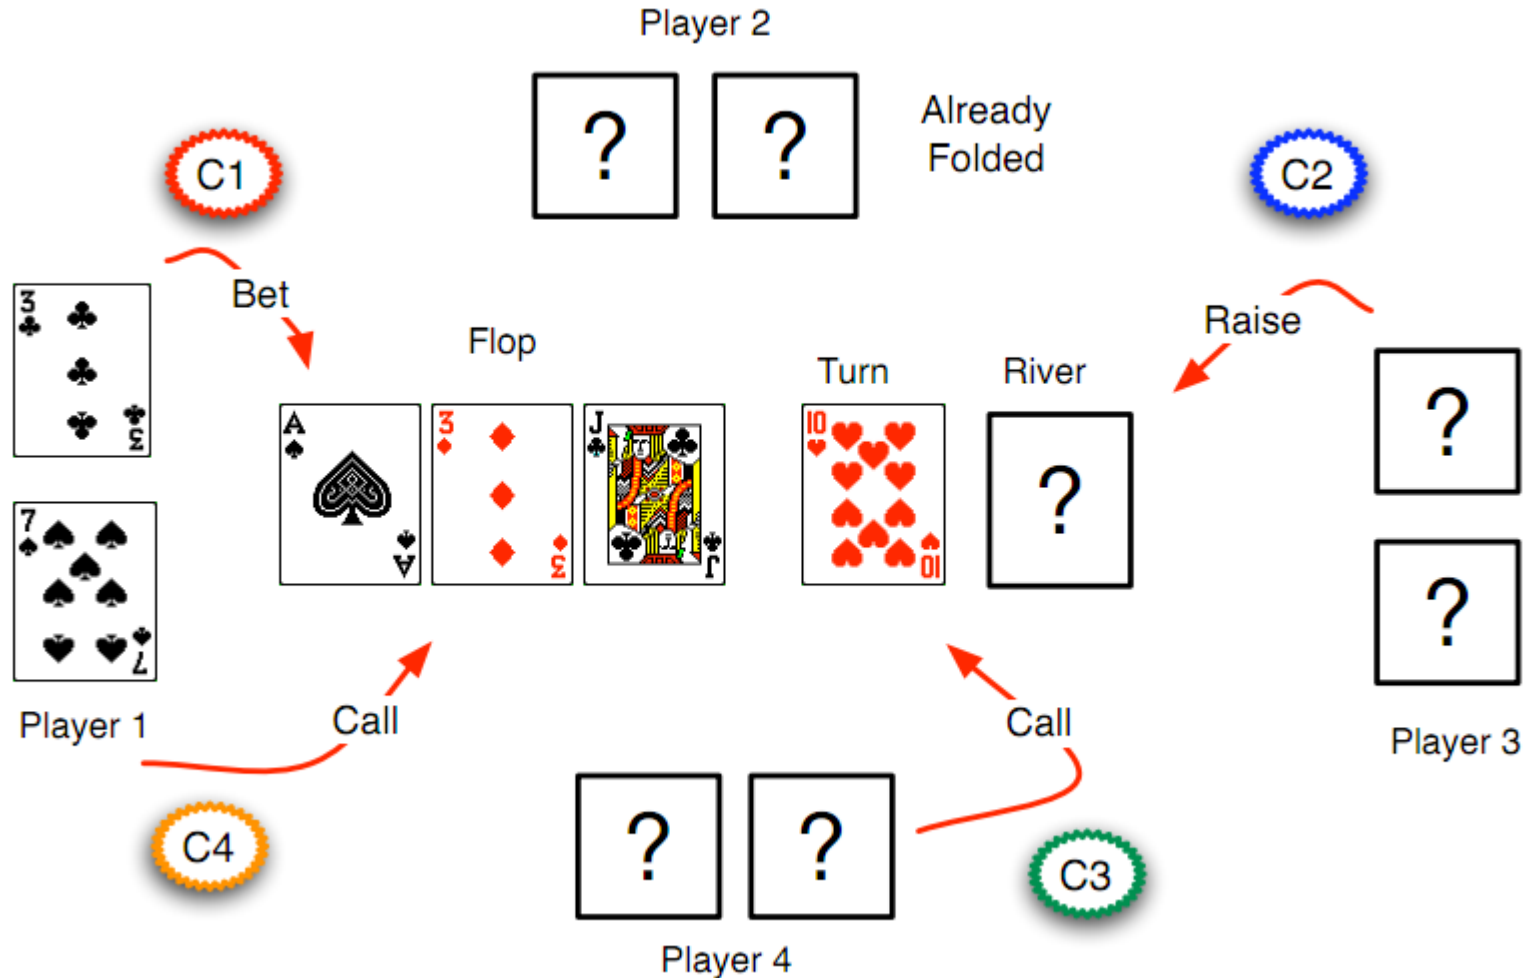
- Call bluffs
- Compare hand strengths

Rule Example


Context: When in the post-flop betting round with 1 raise in the current round, 3 or more players remaining, and pot odds of 0.25. **Action:** player 3 often calls even when his hand is not very good

- **Links Context:**
 - Betting round (post-flop)
 - Nr of raises (1)
 - Nr of players (3 or more)
 - Pot odds (0.25)
- **With Action (call)**
- **And hand strength (not very good)**

Modelling Opponents




Modelling Opponents




Showdown

Player 2



Folded, so
no information
about cards
is known



Player 1


Flop



Turn



River



Player 3



Player 4

Data for the Models

Context	Action	Hole Cards	Shared Cards	Hand Strength
C2	Bet/Raise	$(A\heartsuit, J\spadesuit)$	$(A\spadesuit, 3\heartsuit, J\clubsuit, 10\heartsuit)$	0.950
C6	Bet/Raise	$(A\heartsuit, J\spadesuit)$	$(A\spadesuit, 3\heartsuit, J\clubsuit, 10\heartsuit, 2\heartsuit)$	0.911

Figure 21: Information for the model of player 3.

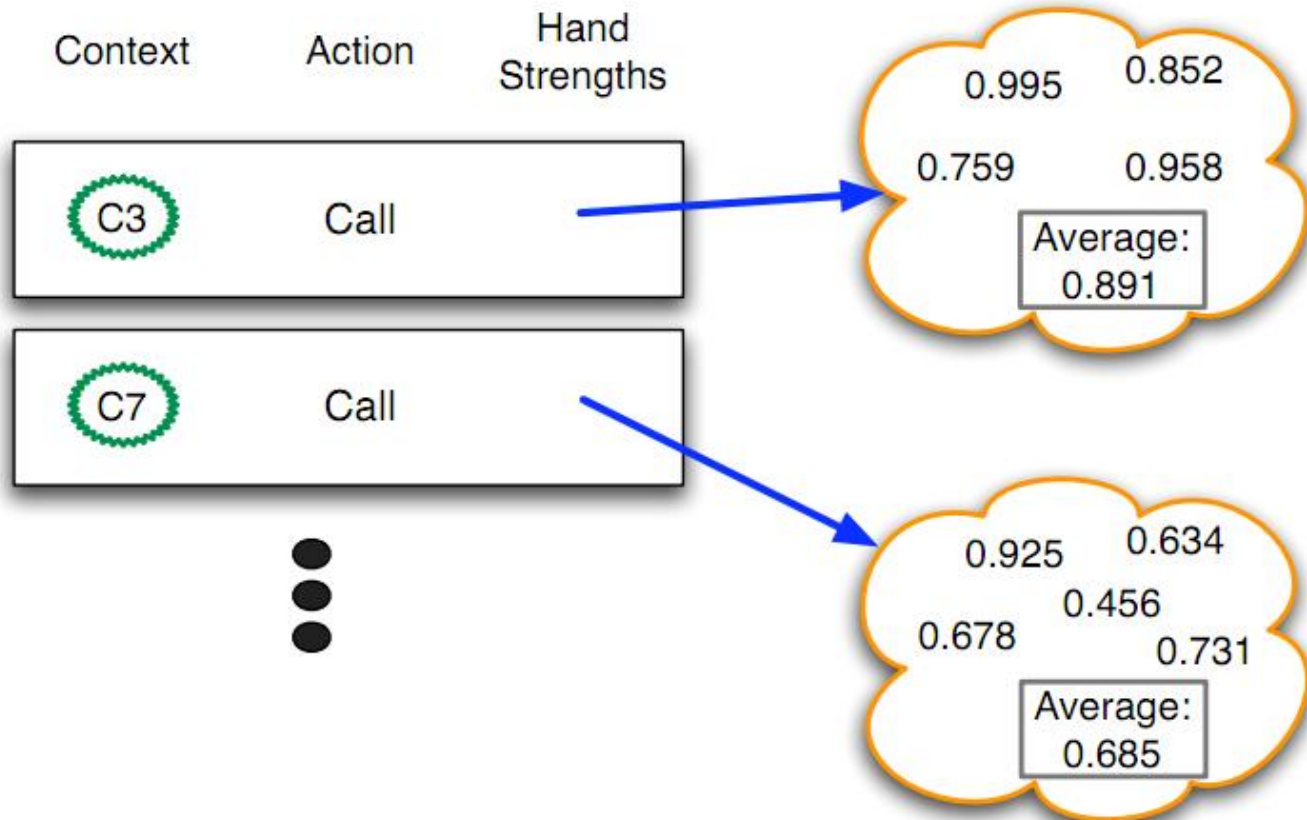
Context	Action	Hole Cards	Shared Cards	Hand Strength
C3	Call	$(10\spadesuit, 10\heartsuit)$	$(A\spadesuit, 3\heartsuit, J\clubsuit, 10\heartsuit)$	0.958
C7	Call	$(10\spadesuit, 10\heartsuit)$	$(A\spadesuit, 3\heartsuit, J\clubsuit, 10\heartsuit, 2\heartsuit)$	0.925



Figure 22: Information for the model of player 4.

$$\text{hand strength} = \left(\frac{\text{Wins} + \frac{\text{Ties}}{2}}{\text{Wins} + \text{Ties} + \text{Losses}} \right)^k$$

Averaging Over Time

Player 4 Model Table





Estimated Hand Strengths

Contexts

Useful features include:

1. The betting round, whether (1) pre-flop, (2) post-flop, (3) post-turn or (4) post-river.
2. The number of players remaining in the hand.
3. The number of raises in the current betting round.
4. The *pot odds* = C:P, where:

C = the amount needed to call the current bet, and

P = the size of the pot (i.e., the total amount bet so far in the current hand).

Expressed as a fraction, the pot odds are:

$$\frac{C}{C + P}$$

Too Specific?

Context

Action

Specialized Situation

Betting Round = 4
Num Players In = 5
Num Raises = 2
 $0.25 \leq \text{Pot Odds} < 0.28$

Call

0.238

Betting Round = 4
Num Raises < 2
Pot Odds > .2

Raise

More General Situation

0.934 0.767
0.552 0.866
0.701

Average:
0.764

Discretising Features

Bin 1: $F \leq 0.1$

Bin 2: $0.1 < F \leq 0.2$

Bin 3: $0.2 < F \leq 0.3$

Bin 4: $F > 0.3$

Is my Context OK?

```
Model of opponent 0 :  
flop, many,few, Low, call, : [0.24752111465140578, 69]  
flop, many,many, Low, raise, : [0.9634597594819612, 2]  
preFlop, many,few, Low, raise, : [0.4612857142857143, 7]  
turn, many,many, High, raise, : [0.5587553190661781, 3]  
river, many,many, Low, call, : [0.7622269095402551, 20]  
river, many,many, Low, raise, : [0.6691929856084127, 70]
```

Tools for Deciding Actions

- Power rating
- Pre-flop winning probability
- Hand strength
- Opponent model

PHASES AND DELIVERABLES



Source: Flickr (by maorix)

URL: <http://www.flickr.com/photos/maor-x/29>

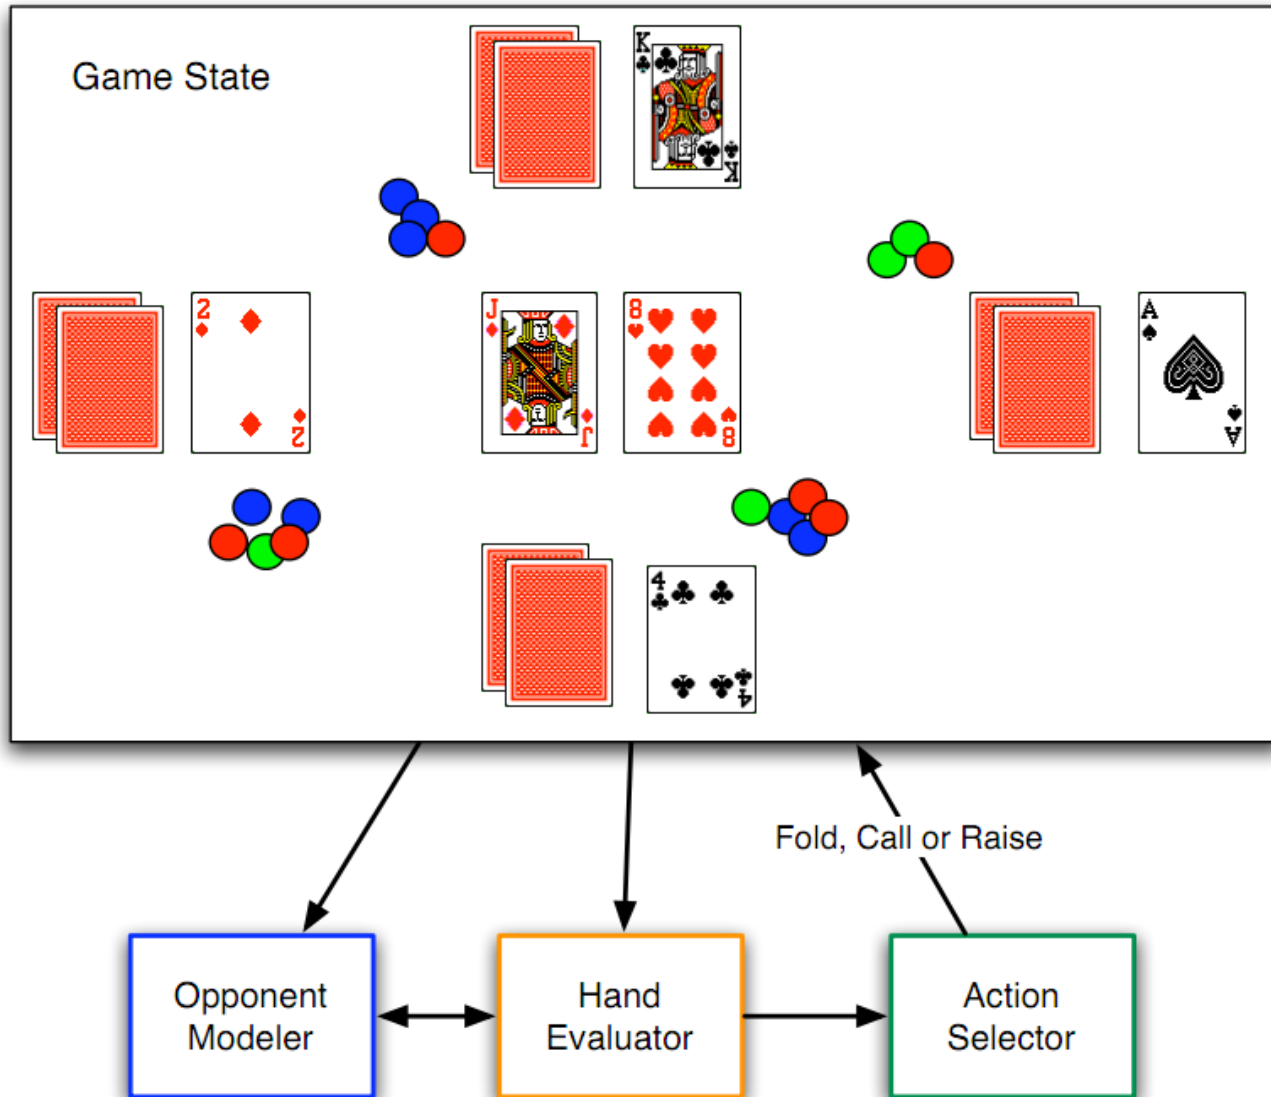
Project Phases

Phase I The basic simulator for k-player Texas Hold 'Em. **25 points.**

Phase II A *pre-flop rollout simulator* used to generate the winning potential of all hole-card pairs, and a procedure for calculating the hand-strength of the combination of hole cards and community cards. **45 points.**

Phase III An *opponent modeler* that a) records associations between game situations, opponent actions and the strengths of their hands, and b) uses those associations to estimate an opponent's hand strength in the current game. **25 points.**

Phase 1: Basic Simulator



Print Some Data!

- Before betting:
 - Hole cards
 - Shared cards
 - Pot
- During betting:
 - Actions
 - Current bet

Hole Cards:

Player 0 : [[10, 'C'], [2, 'C']]

Player 1 : [[13, 'S'], [7, 'S']]

Cash:

Player 0 : 915

Player 1 : 970.0

Common cards: [[9, 'S'], [2, 'S'], [10, 'H'], [14, 'C'], [5, 'H']]

Pot: 115

Player 0 : Hand strength is 0.928282828283 and action is raise

Player 0 does action: raise . Current bet is 10

Player 1 : Hand strength is 0.339898989899 and action is raise

Player 1 does action: raise . Current bet is 20


Player 0 : Hand strength is 0.928282828283 and action is raise

Player 0 does action: raise . Current bet is 30

Player 1 : Hand strength is 0.339898989899 and action is call

Player 1 does action: call . Current bet is 30

Phase 2: Pre-Flop Rollout Simulator



Phase 2: Hand Strength Calculator




Missing Turn and/or River cards are NOT filled in during the basic hand-strength calculation.




- vs -



$\binom{47}{2}$ possible pairs



$$\text{hand strength} = \frac{\text{wins} + (0.5)\text{ties}}{\text{wins} + \text{losses} + \text{ties}}$$

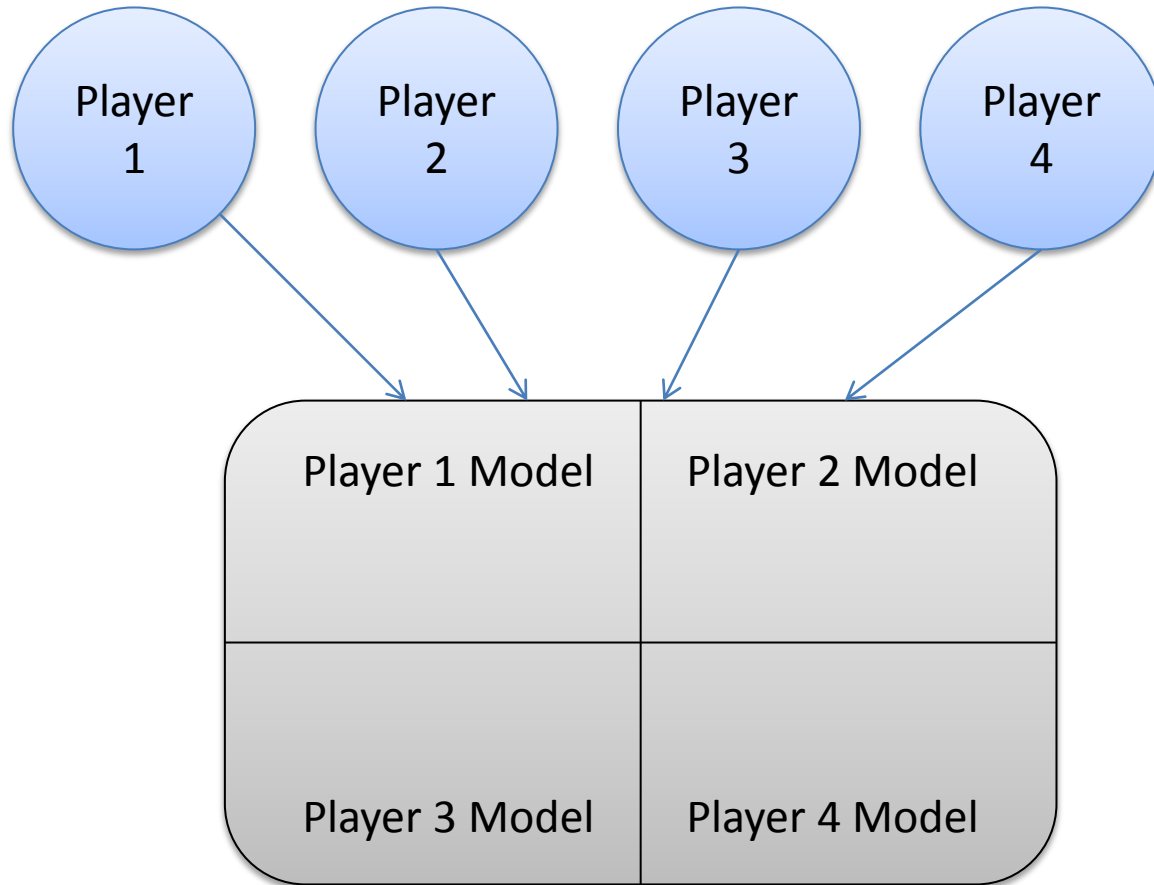


Phase 2: Decisions

- Make wiser decisions by using rollout probabilities and hand strengths.
- Prove they make the player better:


```
1-3: Phase1 4-6: phase2
Player 0 : -13375.0
Player 1 : -11000.0
Player 2 : -9875.0
Player 3 : 33125.0
Player 4 : 32625.0
Player 5 : 28500.0
```

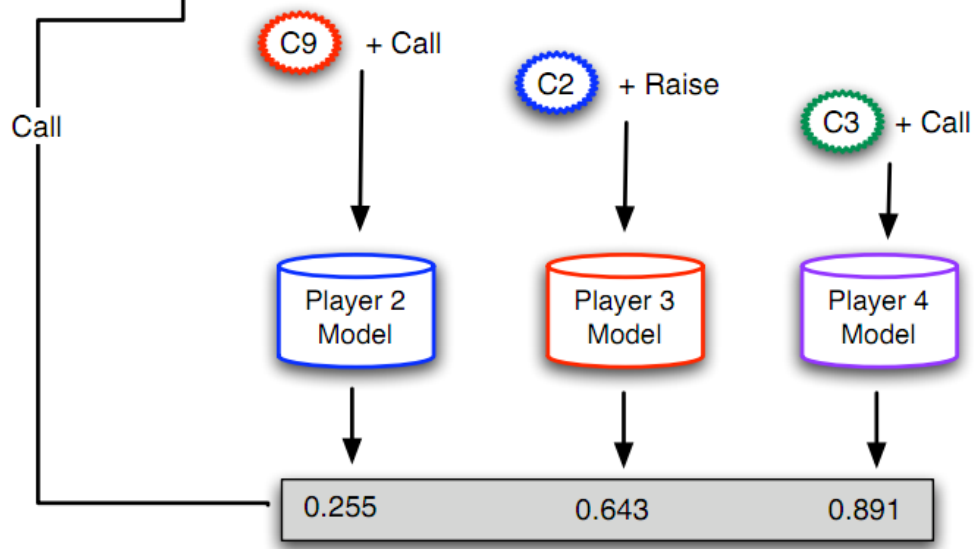
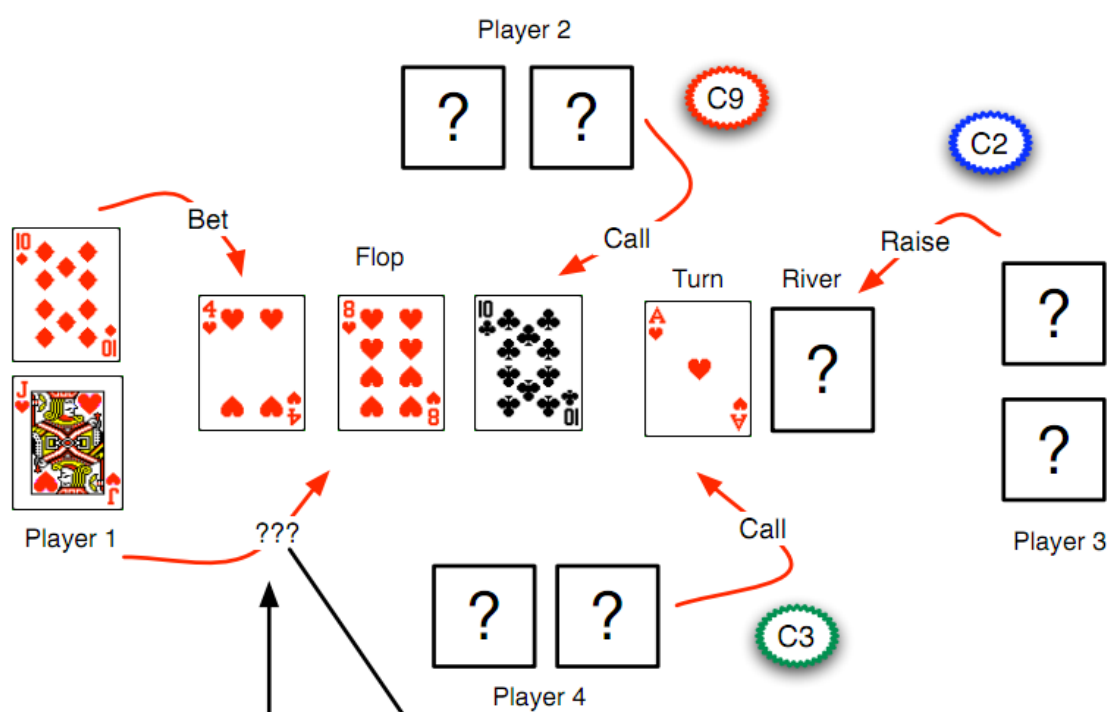
Phase 3: Opponent Modelling



Model: Gathering Data, Using Data

Player 4 Model Table





Estimated Hand Strengths

Using Opponent Data

- Conservative:
if myStrength > bestOpponent: raise
- Liberal:
if myStrength > anyOpponent: raise

Testing Opponent Modeller

Results from:

Player 0 - Conservative opponent modeller (play when better than BEST opponent)

1 & 2 - Quite poor, super-simple.

Player 3 & 4: Phase 2

Player 5: Aggressive opponent modeller (play when better than average.)

Tournament done. Results are:

Player	0	:	51750.0
Player	1	:	-23000.0
Player	2	:	-10625.0
Player	3	:	13000.0
Player	4	:	20000.0
Player	5	:	8875.0

Report

1. The basic structure of your code for each project phase,
2. The logic behind the betting decisions made by the players in **each** of the project phases that you complete,
3. The opponent models used in phase III, in particular, the contexts used to differentiate game situations,
4. The results of multi-hand (1000) runs for **each** of the project phases that you complete. These should **not** be screen dumps, but simple tables indicating each players total winnings. Five simple tables per project phase is all that is required (and desired) here.
5. A brief discussion of why (you believe) certain strategies were more (or less) effective than others for **each** phase of the project.

Grading

1. A well-functioning Phase I player that obeys all fundamental poker rules. **(20 points)**
2. Five different 1000-hand runs of a game involving 4 or more Phase-I players, with the results of each 1000-game segment summarized in a table showing the final resource (e.g. money) of each player. **(5 points)**
3. Modules for computing, storing (to file) and loading (for use in a poker game) the pre-flop rollout probabilities for $k = 2$ to 10 player games. Each rollout for each card equivalence class should involve at least 1000 rounds of simulated play. **(20 points)**
4. Modules for estimating hand strength for a pair of hole cards plus 3, 4 or 5 community/shared cards. This will involve testing the given hand against hands made from all possible pairs of hole cards (using the remaining 47 to 45 cards), as described in ai-poker-players.pdf and diagramed in Figure 15 of that document. **(15 points)**
5. Five different 1000-hand runs of a game involving 4 or more Phase-II players, with the results of each 1000-game segment summarized in a table, as above. **(10 points)**
6. Modules for performing basic opponent modeling. **(15 points)**
7. Five different 1000-hand runs of 4 or more players, at least of 2 of which use Phase-III techniques, while the others may use only Phase-II methods. Include the standard summary table of each 1000-game segment. **(10 points)**
8. The general quality of the report, which must cover all of the 5 topics mentioned above. **(5 points)**.

Some Tips

- Check your models manually!
 - Pre-flop probabilities
 - Hand strengths
 - Opponent models
- Look at game printouts
- Read lecture notes and project description
- Play poker!