Bayesian Poker

Ann E. Nicholson and **Kevin B. Korb**

Faculty of Information Technology
Monash University
Clayton, VIC 3168
AUSTRALIA

{annn,korb}@csse.monash.edu.au
http://www.csse.monash.edu.au/bai

Text: *Bayesian Artificial Intelligence*, Kevin B. Korb and Ann E. Nicholson, Chapman & Hall/CRC, 2004.

Overview

- The Monash Poker Project
- Poker
- A Bayesian Network for poker
- Betting curves for randomisation
- Bluffing
- Opponent Modelling
- Experimentation to date

Poker and AI

- Poker is ideal for testing automated reasoning under uncertainty
  - Physical randomisation
  - Incomplete information
  - Unpredictable opponent (strategies, bluffing, etc)
- Bayesian networks are a good representation for complex game playing.
- Our Bayesian Poker Player (BPP) plays at the level of a good amateur human player.

The History of Monash BPP

Started with Kevin Korb in 1993 (5 card stud).

- Nathalie Jitnah, 1993. Bayesian Poker (the basic BN)
- Scott Thomson, 1994. Bayesian Poker (some improvements)
- Aidan Doyle, 1995. Web interface.

  Int. Conf on Uncertainty in AI, 1999. K.B. Korb, A.E. Nicholson and N. Jitnah (1999) Bayesian Poker.
- Jason Carlton, 2000. Extended to decision networks.
- Darren Boulton, 2002-03. Extended to Texas Hold'em Poker. Improved bluffing and opponent modelling.
- 2006
  - Inauguaral Bot Poker competition, Conference of the American Artificial Intelligence Association. (3rd)
  - New Web Interface (Steven Mascaro), demo at Open Day (this Sunday!).

Five-Card Stud Poker

Players start with an *ante* (initial fixed bet) and are dealt a sequence of 5 cards, first face down (hidden), the rest face up.

Player actions after each card is dealt:

- PASS - first player may make no bet
- CALL - match the current bet on the table
- RAISE - increase the bet
- FOLD - drop out of this hand

Texas Hold'em Poker

Players dealt 2 card face down, then round of betting (pre-flop).

Remaining dealt cards (5) are "shared".

Betting round after 3 shared cards (flop).

Betting round after 4th shared card (turn).

Betting round after 5th shared card (river).

Poker Hands

Poker Hands from weakest to strongest:

Hand Type	Example	Probability	
		5 card stud	Texas Hold'em
Busted	A♣ K♠ J♢ 10♢ 4♡	0.5015629	0.1728400
Pair	2♡ 2♢ J♠ 8♣ 4♡	0.4225703	0.4380000
Two Pair	5♡ 5♣ Q♠ Q♣ K♣	0.0475431	0.2351900
3 of a Kind	7♣ 7♡ 7♠ 3♡ 4♢	0.0211037	0.0483400
Straight	3♠ 4♣ 5♡ 6♢ 7♠	0.0035492	0.0479900
Flush	A♣ K♣ 7♣ 4♣ 2♣	0.0019693	0.0299000
Full House	7♠ 7♢ 7♣ 10♢ 10♣	0.0014405	0.0255000
4 of a Kind	3♡ 3♠ 3♢ 3♣ J♠	0.0002476	0.0018800
Straight Flush	3♠ 4♠ 5♠ 6♠ 7♠	0.0000134	0.0003600

Pot Odds & Probabilities

Let p = probability of winning the pot if the hand is played to showdown,

$n-1$ = number of opponents remaining in game,

k = expected cost of reaching the showdown.

$$\text{pot odds} \quad = \quad \frac{k}{c+k-1} \tag{1}$$

The *calling threshold*, θ, identifies the probability of winning at which the expected values of calling a bet versus folding are equal.

$$\theta \quad = \quad \frac{k}{c+2k-1} \tag{2}$$

Poker BN (circa 99)

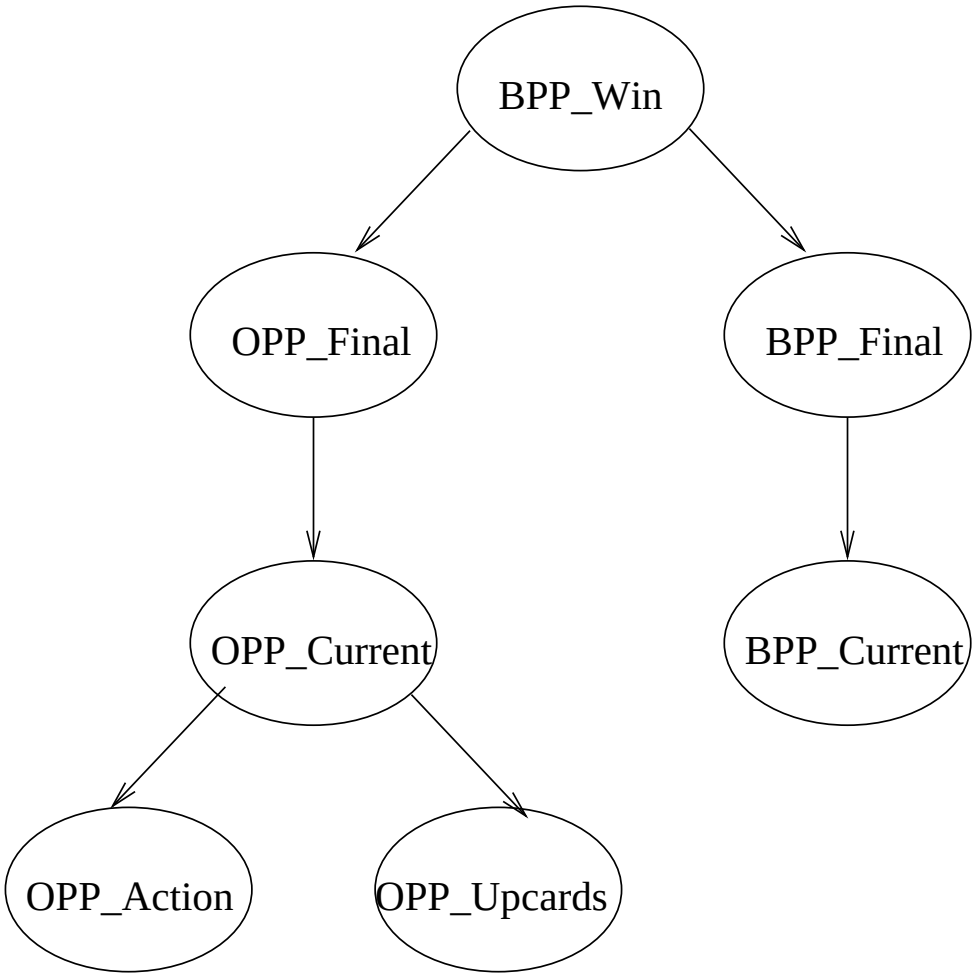

Bayesian Poker Network (cont.)

- 4 different networks: 1 for each round of play.
- *OPP Current*, *BPP Current*: (partial) hand types with cards dealt to now.
- *OPP Final*, *BPP Final*: hand types after all 5 cards dealt.
- Observation nodes:
  - *OPP Upcards*: All opponents cards except first are visible to BPP.
  - *OPP Action*: BPP knows opponents action.

Hand Types

- Initial 9 hand types too coarse.
- Actual number of distinct poker hand values is 7462

  (http://www.suffecool.net/poker/evaluator.html)

Type	Unique	Distinct
Straight Flush	40	10
Four of a Kind	624	156
Full Houses	3744	156
Flush	5108	1277
Straight	10200	10
Three of a Kind	54912	858
Two Pair	123552	858
One Pair	1098240	2860
High Card	1302540	1277
TOTAL	2598960	7462

- Current BPP recognises 25 hand types a pair):
  - busted-low (leq 8) busted-medium (9,10,J) busted-Q-high, busted-K-high, busted-A-high
  - all pairs represented separately
  - 7 other hand-types

Conditional Probability Tables

In general, for BNs:

- Each node has an associated CPT
- CPT gives probability of child node given the combination of values of the parent node.

For this BPP:

- Prior for "Final" hand as per table
- $P(Current|Final)$ obtained by dealing many hands
- $P(OppAction|Current)$ - very simplistic, can start with "rules", learn as play opponent(s)

Belief Updating

Given evidence for *BPP_Current* and *OPP_Action*, belief updating produces

- belief vectors for both players' final hand types
- a posterior probability of BPP winning the game.

A Decision Network for Poker

(Carlton 2000 version)

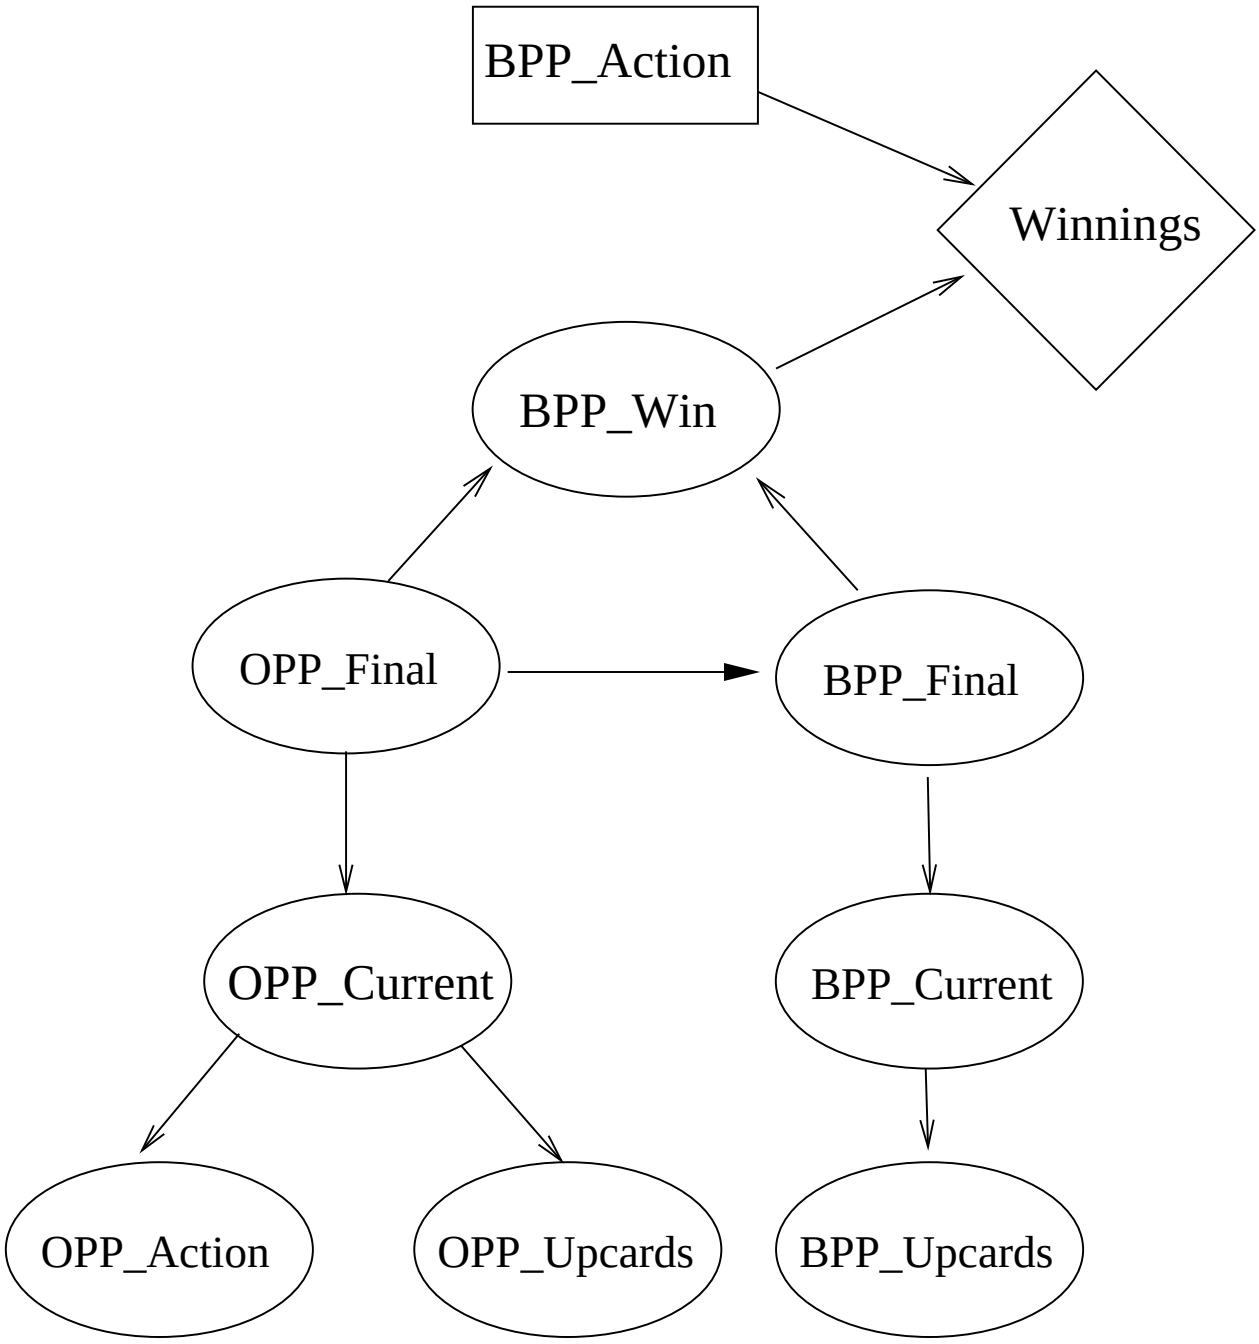

Making the decision

The utility node, *Winnings*, measures the dollar value BPP expects to make based on the possible combinations of the states of the parent nodes (*BPP_Win* and *BPP_Next_Action*).

BPP_Next_Action	*BPP_Win*	Utility
Bet	Win	$F_{bet} - B_{bet}$
Bet	Lose	$-B_{bet}$
Pass	Win	$F_{pass} - B_{pass}$
Pass	Lose	$-B_{pass}$
Fold	Win	0
Fold	Lose	0

B size of betting unit; R no. of betting round remaining; C current pot size.

If betting, need to estimate:

F_{bet} – final pot

B_{bet} – BPP's future contribution

Betting with Randomization

- *Betting curves* are used to randomize betting actions.

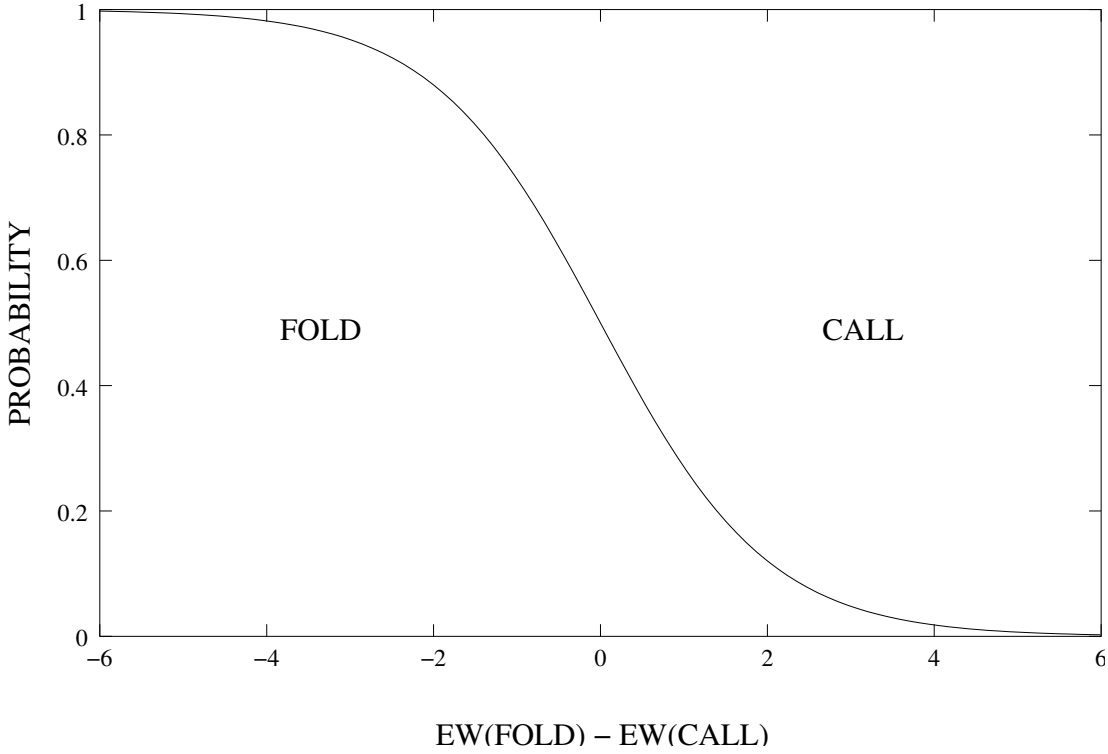

- The horizontal axis shows the difference between the EW of folding and calling.
- Betting curves were generated with exponential functions.
- "Good" parameters found using stochastic search with BPP vs. a Rule-Based Opponent.

Bluffing

- *Bluffing* is the intentional misrepresentation of the strength of one's hand (e.g. over-representing that strength).
- With low probability (5%) BPP will enter "bluffing" mode.
- Bluffing done by modifying predicted belief in winning (by halving probability of losing).
- E.g. If predicted 60% chance of winning, modify to 80%.
- Decisions then made the same way.

Current Decision Network for Texas Hold'em Poker

(Boulton 2003 version)

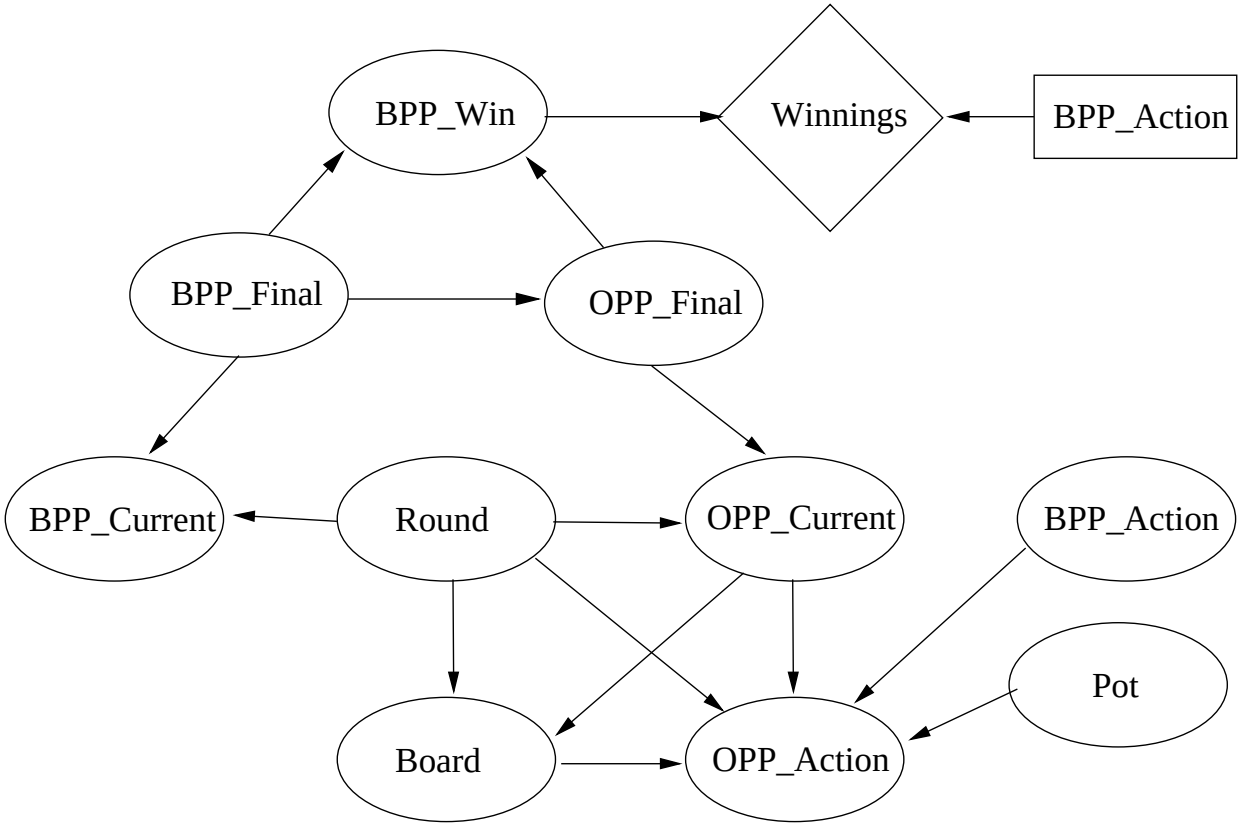

How good is BPP?

- Beats
  - A Simple Probabilistic Opponent
  - A Rule-Based Opponent
  - Some amateur human Opponents
  - Some bots (AAAI competition)
- Loses to
  - Good bots (AAAI competition) - U. of Alberta (Billings group), CMU bot.
  - experienced humans

Extensions

- BN Improvements
  - Adding bluffing to the opponent model
  - Improved learning of opponent model
- More complex poker
  - Multi-opponent games
  - table stake games
- Dynamic Bayesian network to represent play over time