Structure and Synthesis of Robot Motion

Motion Synthesis with Strategic Considerations I

Subramanian Ramamoorthy
School of Informatics

16 February, 2012

Acknowledgement

The contents of this lecture are based on a tutorial at AAAI 2008 by Y. Gal and A. Pfeffer

I have used their original slides, with a bit of editorial additions to fit within a lecture.

Question: How will you tell a humanoid robot to mark another player?

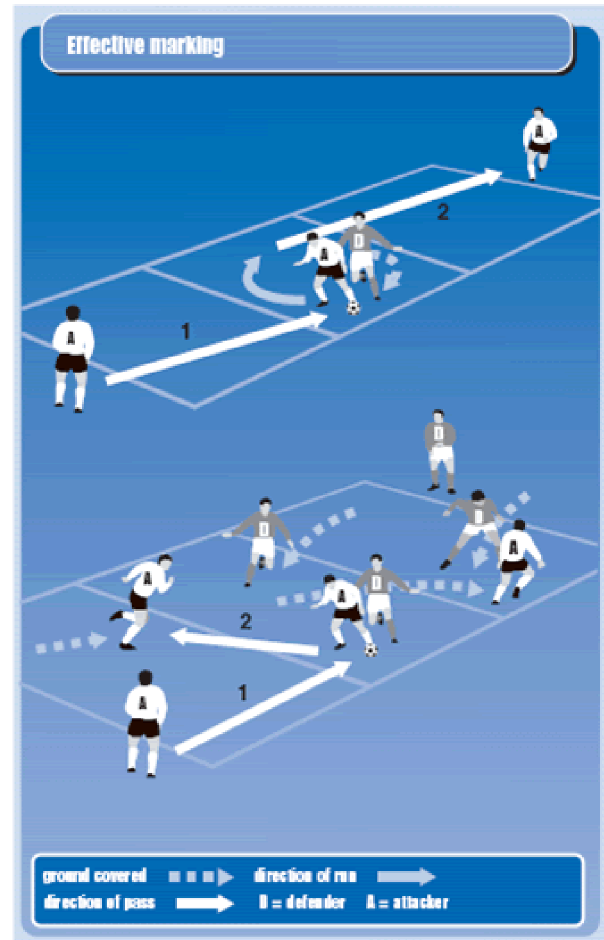

Applications of Societal Importance

NurseBot

Ubiko, a hospital robot guide
Aizu Central Hospital, Japan

Multi-agent Decision-making

- Two sources of uncertainty
  - the environment
  - other agents
- Multi-agent decision problem, or game, includes
  - strategies
  - outcomes
  - utilities
- A solution concept for a game includes a strategy *profile* for all agents.

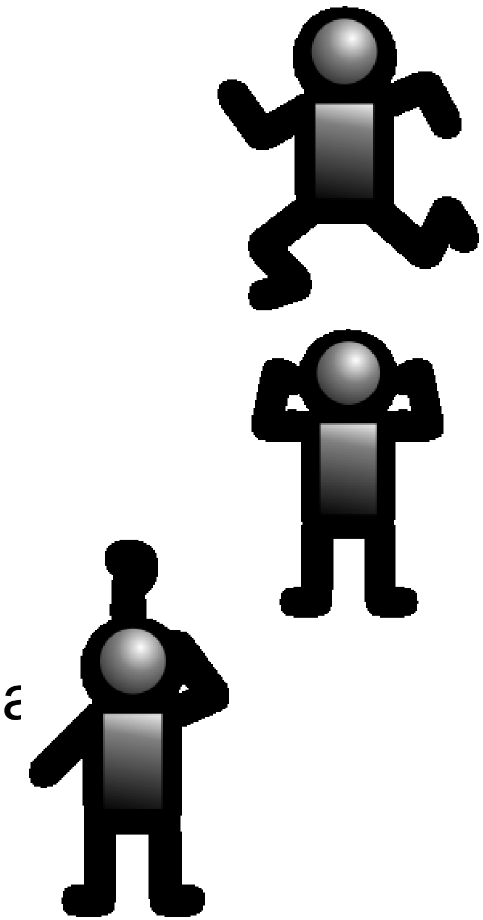

What is a game ?

A game includes a set of agents $\mathbf{N} = \{1, \ldots, n\}$.
For each agent i, includes a set of strategies $\mathbf{S}_i$.

Joint strategy profile $(s_1, \ldots, s_n)$ determines outcome of game, where $s_i \in \mathbf{S}_i$.

Payoff function, $u_i : \mathbf{S}_1, \ldots, \mathbf{S}_n \rightarrow R$
represents utility for i given $(s_1, \ldots, s_n)$

Let $\mathbf{S}_{-i} = \mathbf{S}_1 \times, \ldots, \times \mathbf{S}_{i-1} \times \mathbf{S}_{i+1} \times, \ldots, \times \mathbf{S}_n$ be the joint set of strategies for all players other than i.

Normal Form Representation: The Prisoners' Dilemma

		Bob	
		C	*D*
Alice	*C*	*(-1,-1)*	*(-5,0)*
	D	*(0,-5)*	*(-3,-3)*

- Row player is Alice; column player is Bob. Values in cell *(C,D)* denotes payoff to Alice when playing *C* and to Bob when playing *D.*
- A **dominant** strategy is one which is best for an agent regardless of other agents' actions.
- The dominant strategy for both players in the prisoners' dilemma is to defect *(D,D)*.

Battle of the Sexes

- Row player is Alice; column player is Bob.
  - Alice prefers watching a football match (*FM*) over going to the ballet (*B*); conversely for Bob.
  - Both players do not like to mis-coordinate.
- Entry *(FM, B)* denotes payoff to (Alice, Bob) when Alice goes to *FM* and Bob goes to *B*.

Bob

		FM	*B*
Alice	*FM*	*(2,1)*	*(0,0)*
	B	*(0,0)*	*(1,2)*

Battle of the Sexes

- Dominant strategies do not exist for either Alice or Bob.
  - But given Alice's strategy, Bob can choose a strategy to maximize his utility (and similarly for Alice)

Bob

Alice		*FM*	*B*
	FM	*(2,1)*	*(0,0)*
	B	*(0,0)*	*(1,2)*

Battle of the Sexes

- Dominant strategies do not exist for either Alice or Bob.
  - But given Alice's strategy, Bob can choose a strategy to maximize his utility (and similarly for Alice)

Bob

Alice

	FM	*B*
FM	***(2,1)***	*(0,0)*
B	*(0,0)*	***(1,2)***

Nash Equilibrium

- A strategy profile $\mathbf{s}^* = (s_1, \ldots, s_n)$ is a Nash equilibrium if no player has the incentive to deviate from its assigned strategy.

- Formally, for every player i and $s_i \in \mathbf{S}_i$,

$$u_i(s_i^*, \mathbf{s}_{-i}^*) \geq u_i(s_i, \mathbf{s}_{-i}^*)$$

Bob

Alice		*FM*	*B*
	FM	*(2,1)*	*(0,0)*
	B	*(0,0)*	*(1,2)*

Matching Pennies

- Alice and Bob can each turn a penny to "heads" or "tails". The payoffs depend on whether Alice and Bob coordinate. The game is zero sum.

Bob

Alice		*H*	*T*
	H	*(1,-1)*	*(-1,1)*
	T	*(-1,1)*	*(1,-1)*

Mixed Strategies

For each player i, a mixed strategy profile defines a probability $\sigma_i(s_i)$ for each pure strategy s_i.
Let σ be a mixed strategy profile for all players.

The expected utility for i given σ is

$$u_i(\sigma) = \sum_{\mathbf{s} \in \mathbf{S}} \prod_j \sigma_j(s_j) \cdot u_i(\mathbf{s})$$

Mixed Strategy Equilibrium

σ is a Nash equilibrium if no player has the incentive to deviate from its assigned strategy.

Formally, for every player i and $\sigma_i \in \Delta \mathbf{S}_i$,

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geq u_i(\sigma_i, \sigma_{-i}^*)$$

Theorem (Nash 50) Every finite game has a mixed strategy equilibrium.

Mixed strategy equilibrium for Matching Pennies: Alice and Bob choose heads and tails with probability 0.5.

Sequential Decisions

- Normal form represents situation where agents make simultaneous decisions
- What happens when players make decisions sequentially?
- We need to be able to represent situations in which different agents have different information

- **Extensive form games:** like single agent decision trees, plus information sets

“Constrained” Poker [Kuhn 1950]

- Two players (P1,P2) each given £2
- Three cards in the deck: K, Q, J
- All players put £1 in the pot and pick a card, visible only to themselves.
- P1 bets £1 or passes;
- P2 bets £1 or passes;
- if P1 passes and P2 bets
- P1 can bet its £1 or pass.
- If both players bet (or pass), player with higher card wins £2 (or £1).
- If one player passes and the other bets, the betting player wins £1

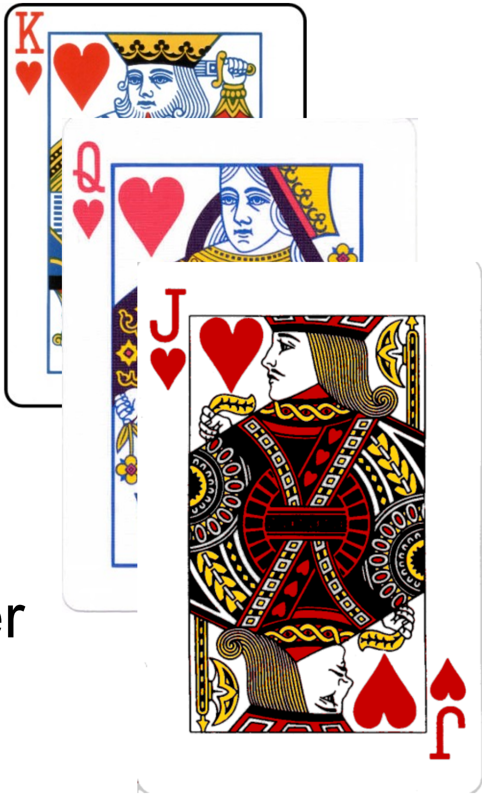

Partial Extensive Form Game Tree for Kuhn Poker

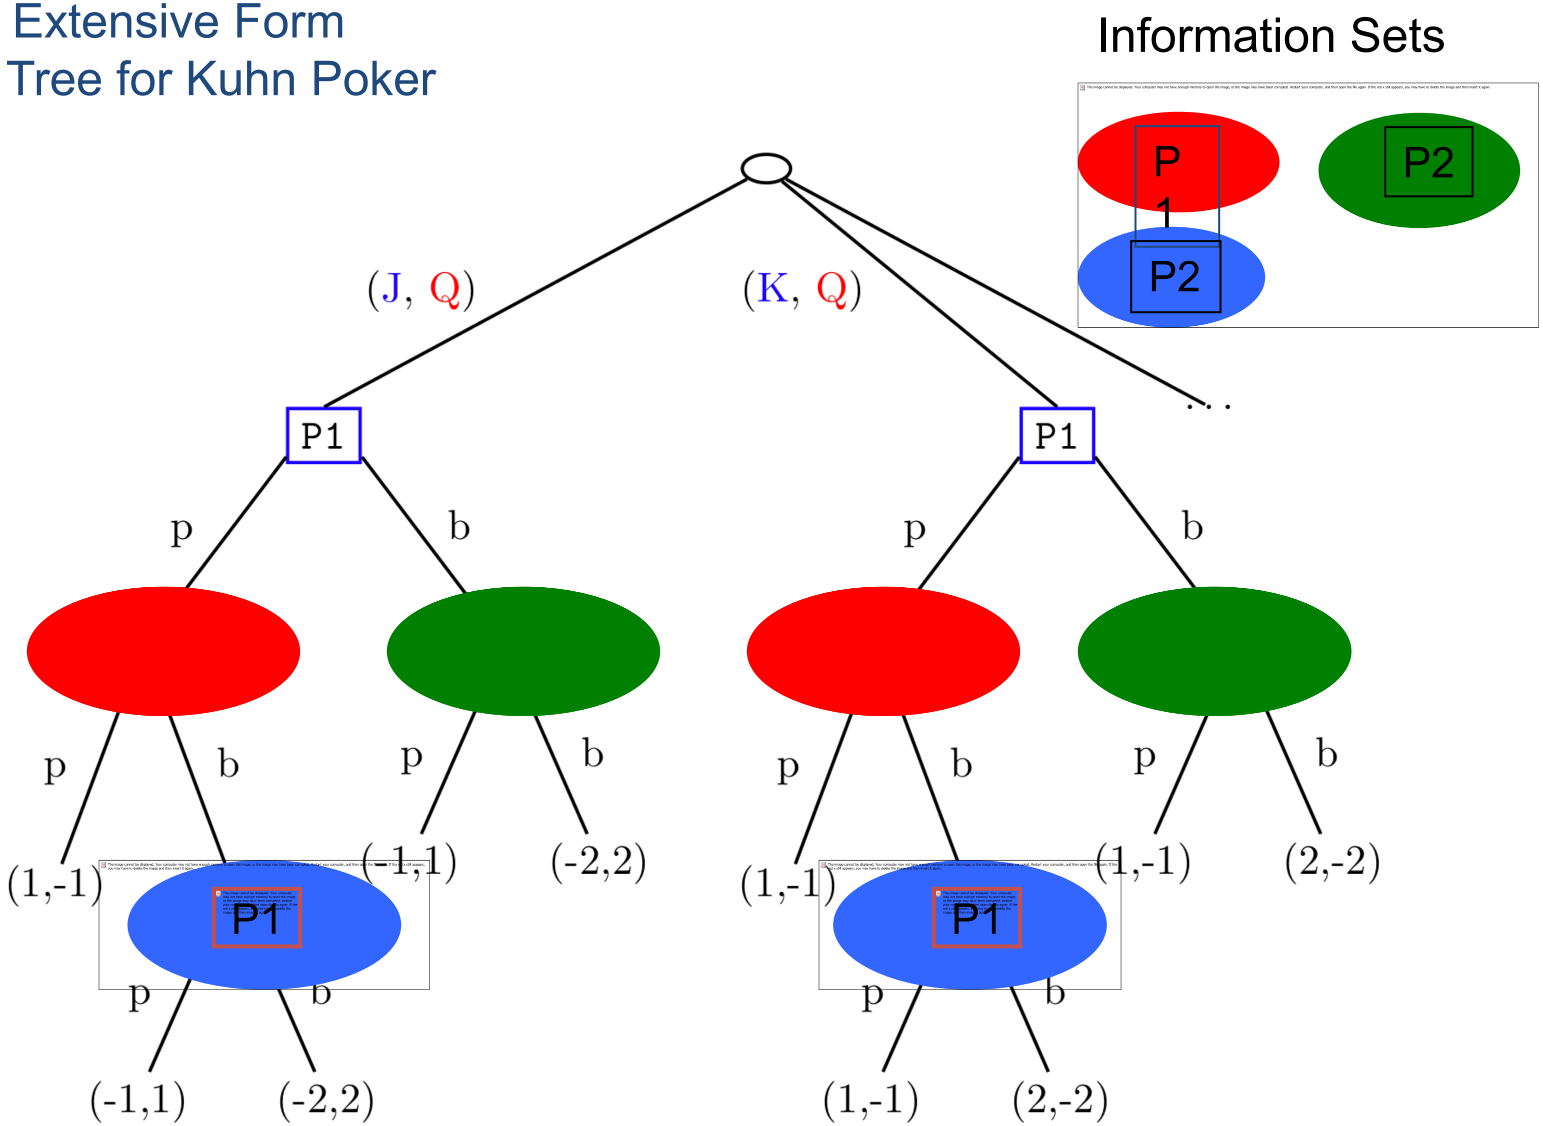

Solution Algorithms: Normal Form Game

- Exact solutions:
- Two player zero-sum games
  - Can be solved by a linear program in polynomial time (in number of strategies)
- Two player general-sum game
  - Can be solved by a linear complimentary program (exponential worst-case complexity) [Lemke-Howson '64]
- Approximate solutions for multi-player games:
  - Continuation and triangulation methods [Govindan and Wilson '03, McKelvey & McLennan '96]
  - Search [Porter, Nudelman and Shoham '05]
- Off-the-shelf packages

  Gambit, Game tracer

Solution Algorithms: Extensive Form Games

- Two player perfect information zero-sum game
  - can be solved by minimax search (with alpha-beta pruning)

- Two player perfect information general-sum game
  - Can be solved using backward-induction

- Two player imperfect information general-sum game
  - Can be solved using sequence form algorithm

    [Koller, Megiddo, von Stengel ‘94].