

SEARCH METHODS IN AI

GAME TREE SEARCH



Arijit Mondal & Partha P Chakrabarti
Indian Institute of Technology Kharagpur

HEURISTIC SEARCH

❑ STATE or CONFIGURATION:

- A set of variables which define a state or configuration
- Domains for every variable and constraints among variables to define a valid configuration

❑ STATE TRANSFORMATION RULES or MOVES:

- A set of RULES which define which are the valid set of NEXT STATE of a given State
- It also indicates who can make these Moves (OR Nodes, AND nodes, etc)

❑ STATE SPACE or IMPLICIT GRAPH

- The Complete Graph produced out of the State Transformation Rules.
- Typically too large to store. Could be Infinite.

❑ INITIAL or START STATE(s), GOAL STATE(s)

❑ SOLUTION(s), COSTS

- Depending on the problem formulation, it can be a PATH from Start to Goal or a Sub-graph of And-ed Nodes


❑ HEURISTICS

- Estimates of cost from a given state to goal. This, along with the current cost of the path from start till now is used to guide the search

❑ HEURISTIC SEARCH ALGORITHMS

- Algorithm A*, Depth-First Branch & Bound, IDA*, AO*, Alpha-Beta, etc
- Knowledge vs Search

GAMES




MULTI-PLAYER GAMES



South Deals
N-S Vul

				♠ K J 8 5	
				♥ J 10 2	
				♦ J 8	
				♣ A J 10 4	
♠ 10 6					♠ Q 9 2
♥ A 6 4 3					♥ 8 7 5
♦ A K 10 5 2					♦ Q 9 6 3
♣ Q 2					♣ K 8 6
				♠ A 7 4 3	
				♥ K Q 9	
				♦ 7 4	
				♣ 9 7 5 3	
<i>West</i>	<i>North</i>	<i>East</i>	<i>South</i>		
1 ♦	Dbl	2 ♦	Pass		
3 ♦	3 ♠	All pass	2 ♠		

PROBABILISTIC GAMES



ROBOT GAMES



PRISONER'S DILEMMA

- Two members of a criminal gang are arrested and imprisoned. Each prisoner is in solitary confinement with no means of communicating with the other.
- The prosecutors lack sufficient evidence to convict the pair on the principal charge, but they have enough to convict both on a lesser charge.
- Simultaneously, the prosecutors offer each prisoner a bargain.

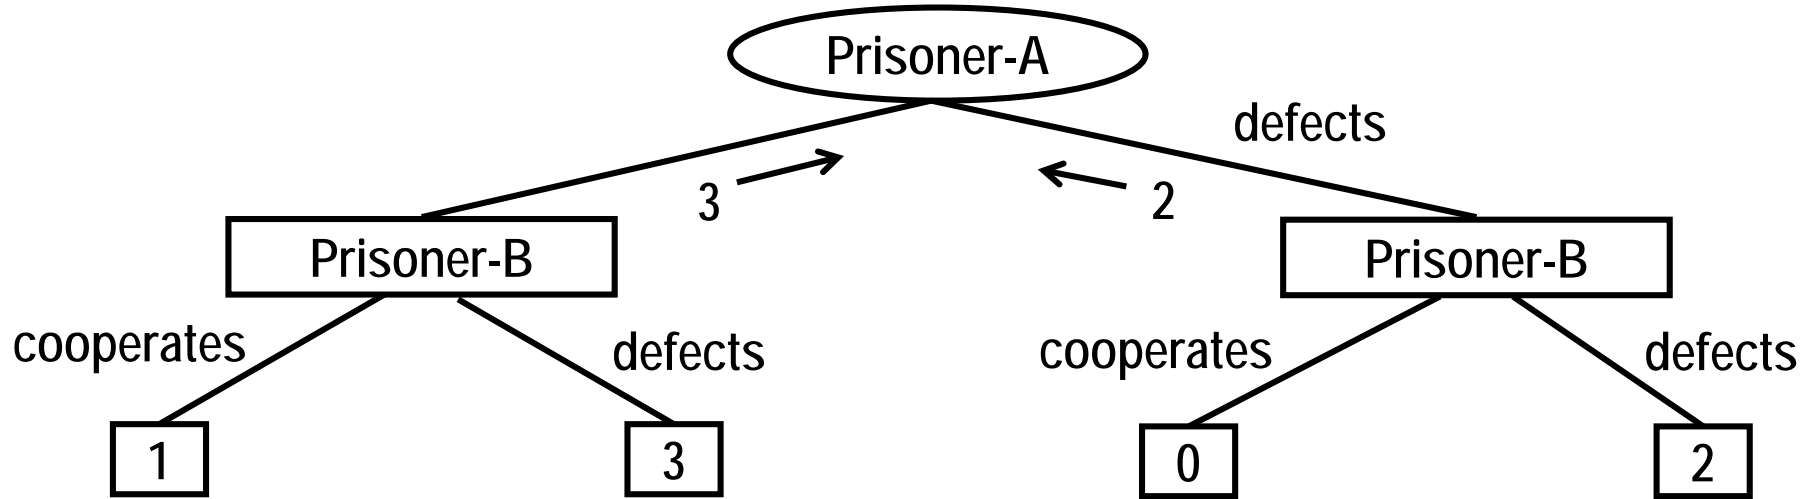
		PRISONER B	
		Prisoner B stays silent (cooperates)	Prisoner B betrays (defects)
PRISONER A	Prisoner A stays silent (cooperates)	Each serves 1 year	Prisoner A: 3 yrs Prisoner B: goes free
	Prisoner A betrays (defects)	Prisoner A: goes free Prisoner B: 3 yrs	Each serves 2 yrs

PRISONER'S DILEMMA: SOLUTION ARGUMENT

- Prisoner A will defect. Why?
 - Prisoner-A knows that Prisoner-B may cooperate or defect
 - If Prisoner-B cooperates, then by defecting, Prisoner-A will go free
 - If Prisoner-B defects, then Prisoner-A will face a longer sentence by staying silent
 - In both cases Prisoner-A gains by defecting
 - Therefore both prisoners will defect, although they would have gained from cooperating



		PRISONER B	
		Prisoner B stays silent (cooperates)	Prisoner B betrays (defects)
PRISONER A	Prisoner A stays silent (cooperates)	Each serves 1 year	Prisoner A: 3 yrs Prisoner B: goes free
	Prisoner A betrays (defects)	Prisoner A: goes free Prisoner B: 3 yrs	Each serves 2 yrs

PRISONER'S DILEMMA: GAME TREE



		PRISONER B	
		Prisoner B stays silent (cooperates)	Prisoner B betrays (defects)
PRISONER A	Prisoner A stays silent (cooperates)	Each serves 1 year	Prisoner A: 3 yrs Prisoner B: goes free
	Prisoner A betrays (defects)	Prisoner A: goes free Prisoner B: 3 yrs	Each serves 2 yrs

TIC-TAC-TOE, CHESS




GAME TREES

- A tree with three types of nodes, namely Terminal nodes, Min nodes and Max nodes. Terminal nodes have no children. The tree has alternating levels of Max and Min nodes, representing the turns of Player-1 and Player-2 in making moves
- All nodes represent some state of the game
- Terminal nodes are labeled with the payoff for Player-1. It could be Boolean (such as WON or LOST). In large games, where looking ahead up to the WON / LOST states is not feasible, the payoff at a terminal node may represent a heuristic cost representing the quality of the state of the game from Player-1's perspective
- The payoff at a Min node is the minimum among the payoffs of its successors
- The payoff at a Max node is the maximum among the payoffs of its successors
- If Player-1 aims to maximize its payoff, then it represents Max nodes, else it represents Min nodes.


SAMPLE GAME TREE

□ Max node : Player-1
○ Min node: Player-2




GAME TREE: MINMAX VALUE

□ Max node : Player-1
○ Min node: Player-2




GAME TREE: DEPTH-FIRST SEARCH (Pruning)

□ Max node : Player-1
○ Min node: Player-2




GAME TREE: DEPTH-FIRST SEARCH (Reverse)




PRUNING IN GAME TREE SEARCH

□ Max node : Player-1
○ Min node: Player-2




PRUNING RULES IN GAME TREE SEARCH



Shallow Cut-off

Exploring node C and its successors is meaningless because payoff at B is at least 14, and hence A will never choose the move to B.



Deep Cut-off

Exploring node E and its successors is meaningless because the payoff at node D is at most 5, whereas the ROOT can already guarantee a payoff of 10 by choosing the left move. Therefore, the game will never reach the node D.

Applicable in a similar manner for a Max node with Min Parent or Ancestor

ALPHA-BETA PRUNING IN GAME TREE SEARCH

- Alpha Bound of J (α) :
 - The max current payoff of all MAX ancestors of J (**Lower Bound**)
 - Exploration of a min node, J, is stopped when its payoff β (**Upper Bound**) equals or falls below alpha.
- Beta Bound of J (β):
 - The min current payoff of all MIN ancestors of J (**Upper Bound**)
 - Exploration of a max node, J, is stopped when its payoff α (**Lower Bound**) equals or exceeds beta
- In a max node, we update alpha or **Lower Bound**
- In a min node, we update beta or **Upper Bound**
- **In both min and max nodes, we return when $\alpha \geq \beta$**


ALPHA-BETA PRUNING PROCEDURE $V(J; \alpha, \beta)$

1. If J is a terminal, return $V(J) = h(J)$.
2. If J is a max node:
 - For each successor J_k of J in succession:
 - Set $\alpha = \max \{ \alpha, V(J_k; \alpha, \beta) \}$
 - If $\alpha \geq \beta$ then return β , else continue
 - Return α
3. If J is a min node:
 - For each successor J_k of J in succession:
 - Set $\beta = \min \{ \beta, V(J_k; \alpha, \beta) \}$
 - If $\alpha \geq \beta$ then return α , else continue
 - Return β

The initial call is with
 $V(\text{Root}; -\infty, +\infty)$

ALPHA-BETA PRUNING PROCEDURE $V(J; \alpha, \beta)$

The initial call is with
 $V(\text{Root}; -\infty, +\infty)$



OTHER ASPECTS IN GAME TREE SEARCH

- Incorporate Heuristics in Game Trees
- Perform Best First Search in Game Trees
- Multi-Player Games for more than two players
- Team Games – Cooperation and Competition
- Probabilistic Games
- Real Life Situations
 - Economics
 - Reactive Control Systems
 - Autonomous Systems

OTHER ASPECTS IN GAME TREE SEARCH

- Incorporate Heuristics in Game Trees
- Perform Best First Search in Game Trees
- Multi-Player Games for more than two players
- Team Games – Cooperation and Competition
- Probabilistic Games
- Real Life Situations
 - Economics
 - Reactive Control Systems
 - Autonomous Systems

OTHER ASPECTS IN GAME TREE SEARCH

- Incorporate Heuristics in Game Trees
- Perform Best First Search in Game Trees
- Multi-Player Games for more than two players
- Team Games – Cooperation and Competition
- Probabilistic Games
- Real Life Situations
 - Economics
 - Reactive Control Systems
 - Autonomous Systems

Thank you