

**CS 1571 Introduction to AI**  
**Lecture 10**

**Finding optimal configurations**  
**Adversarial search**

**Milos Hauskrecht**  
[milos@cs.pitt.edu](mailto:milos@cs.pitt.edu)  
5329 Sennott Square

---

CS 1571 Intro to AI

M. Hauskrecht

**Announcements**

- **Homework assignment 4 is out**
  - **Due on Thursday next week !!!!**
  - Programming and experiments
  - Simulated annealing + Genetic algorithm
  - Competition

**Course web page:**

<http://www.cs.pitt.edu/~milos/courses/cs1571/>

---

CS 1571 Intro to AI

M. Hauskrecht

## Parametric optimization

### Optimal configuration search:

- Configurations are described in terms of variables and their values
- Each configuration has a quality measure  $f$
- Goal: find the configuration with the best value of  $f$

When the state space we search is finite, the search problem is called a **combinatorial optimization problem**

When parameters we want to find are real-valued  
– **parametric optimization problem**

## Parametric optimization

### Parametric optimization:

- Configurations are described by a vector of variables (free parameters)  $\mathbf{w}$  with *real* values
- **Goal:** find the set of parameters  $\mathbf{w}$  that optimize the quality measure  $f(\mathbf{w})$

## Parametric optimization techniques

- **Special cases (with efficient solutions):**
  - Linear programming
  - Quadratic programming
  - Convex programming
- **First-order methods:**
  - Gradient-ascent (descent)
  - Conjugate gradient
- **Second-order methods:**
  - Newton-Rhaphson methods
  - Levenberg-Marquardt
- **Constrained optimization:**
  - Lagrange multipliers

CS 1571 Intro to AI

M. Hauskrecht

## Linear programming

- **A special case and when:**
  - The objective function  $f$  is a linear combination of variable values  $w$
  - Values variables  $w$  can take are constrained by a set of linear constraints
- **Assume variables:**  $w_1, w_2, \dots, w_k$

$$\text{Minimize}_{w_1, w_2, \dots, w_k} f(w_1, w_2, \dots, w_k) = a_1 w_1 + a_2 w_2 + \dots + a_k w_k$$

**Subject to constraints:**

$$b_{1,1} w_1 + b_{1,2} w_2 + \dots + b_{1,k} w_k + b_{1,0} \leq 0$$

$$b_{2,1} w_1 + b_{2,2} w_2 + \dots + b_{2,k} w_k + b_{2,0} \leq 0$$

...


$$b_{m,1} w_1 + b_{m,2} w_2 + \dots + b_{m,k} w_k + b_{m,0} \leq 0$$

CS 1571 Intro to AI

M. Hauskrecht

## Gradient ascent method

- A method for finding parameters  $\mathbf{w}=[w_1, w_2, \dots, w_k]$  optimizing an arbitrary differentiable function  $f(w_1, w_2, \dots, w_k)$
- **Example:**




CS 1571 Intro to AI

M. Hauskrecht

## Gradient ascent method

- **Gradient ascent:** the same as hill-climbing, but in the continuous parametric space  $\mathbf{w}$




- What is the derivative of an increasing function?

CS 1571 Intro to AI

M. Hauskrecht

## Gradient ascent method

- **Gradient ascent:** the same as hill-climbing, but in the continuous parametric space  $w$




- What is the derivative of an increasing function?
  - positive

CS 1571 Intro to AI

M. Hauskrecht

## Gradient ascent method

- **Gradient ascent:** the same as hill-climbing, but in the continuous parametric space  $w$




- Change the parameter value of  $w$  according to the gradient

$$w \leftarrow w^* + \alpha \frac{\partial}{\partial w} f(w) |_{w^*}$$

CS 1571 Intro to AI

M. Hauskrecht

## Gradient ascent method



- New value of the parameter

$$w \leftarrow w^* + \alpha \frac{\partial}{\partial w} f(w) |_{w^*}$$


$\alpha > 0$  - a learning rate (scales the gradient changes)

CS 1571 Intro to AI

M. Hauskrecht

## Gradient ascent method

- To get to the function minimum repeat (iterate) the gradient based update few times



- **Problems:** local optima, saddle points, slow convergence
- More complex optimization techniques use additional information (e.g. second derivatives)

CS 1571 Intro to AI

M. Hauskrecht

## Adversarial search

## Game search

- Game-playing programs developed by AI researchers since the beginning of the modern AI era
  - Programs playing chess, checkers, etc (1950s)
- **Specifics of the game search:**
  - Sequences of player's decisions **we control**
  - Decisions of other player(s) **we do not control**
- **Contingency problem:** many possible opponent's moves must be "covered" by the solution  
Opponent's behavior introduces an uncertainty in to the game
  - We do not know exactly what the response is going to be
- **Rational opponent** – maximizes its own **utility (payoff) function**

## Types of game problems

- Types of game problems:
  - **Adversarial games:**
    - win of one player is a loss of the other
  - **Cooperative games:**
    - players have common interests and utility function
  - **A spectrum of game problems in between the two:**

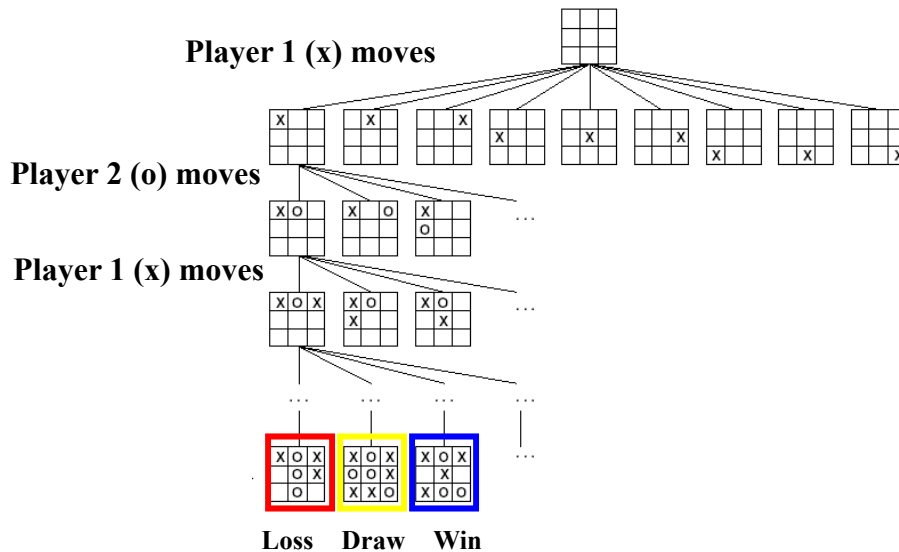
Adversarial games

Fully cooperative games



**we focus on adversarial games only!!**

## Example of an adversarial 2 person game: Tic-tac-toe





## Game search problem

- **Game problem formulation:**
  - **Initial state:** initial board position + info whose move it is
  - **Operators:** legal moves a player can make
  - **Goal (terminal test):** determines when the game is over
  - **Utility (payoff) function:** measures the outcome of the game and its desirability
- **Search objective:**
  - find the sequence of player's decisions (moves) maximizing its utility (payoff)

**Caveat:** Consider the opponent's moves and their utility

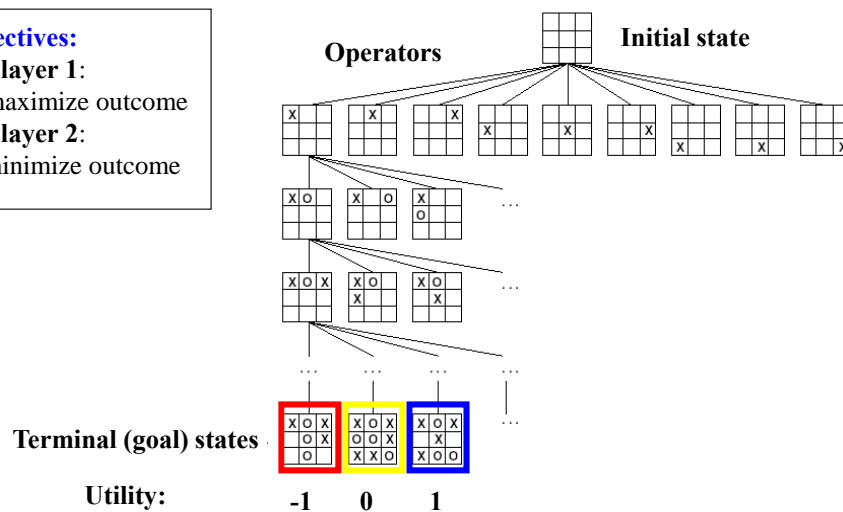
CS 1571 Intro to AI

M. Hauskrecht

## Game problem formulation (Tic-tac-toe)

### Objectives:

- **Player 1:** maximize outcome
- **Player 2:** minimize outcome




CS 1571 Intro to AI

M. Hauskrecht

## Minimax algorithm

### How to deal with the contingency problem?


- Assuming that the opponent is rational and always optimizes its behavior (opposite to us) we consider **the best opponent's response**
- Then the **minimax algorithm** determines the best move



CS 1571 Intro to AI

M. Hauskrecht


## Minimax algorithm. Example



CS 1571 Intro to AI

M. Hauskrecht

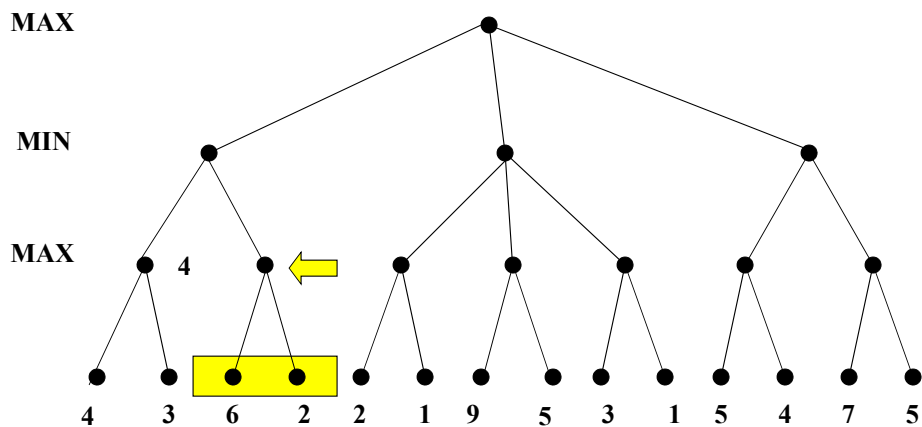
### Minimax algorithm. Example



CS 1571 Intro to AI

M. Hauskrecht


### Minimax algorithm. Example



CS 1571 Intro to AI

M. Hauskrecht

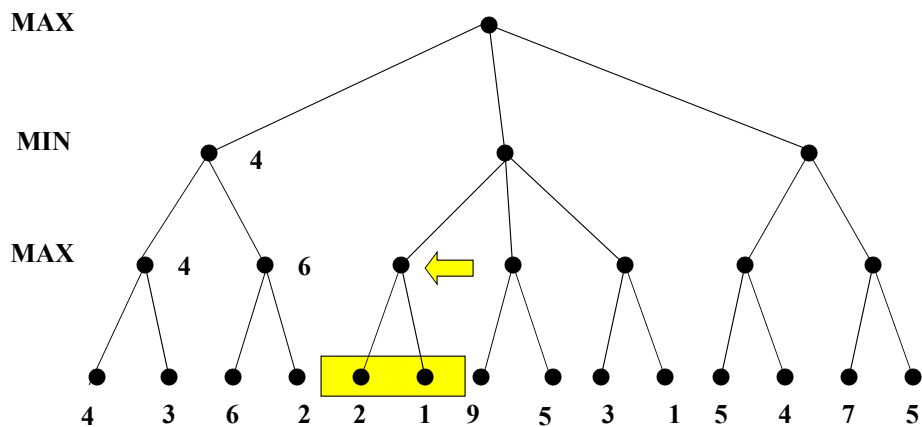
## Minimax algorithm. Example



CS 1571 Intro to AI

M. Hauskrecht


## Minimax algorithm. Example



CS 1571 Intro to AI

M. Hauskrecht


## Minimax algorithm. Example



CS 1571 Intro to AI

M. Hauskrecht


## Minimax algorithm. Example



CS 1571 Intro to AI

M. Hauskrecht


## Minimax algorithm. Example



CS 1571 Intro to AI

M. Hauskrecht


## Minimax algorithm. Example



CS 1571 Intro to AI

M. Hauskrecht

## Minimax algorithm. Example



CS 1571 Intro to AI

M. Hauskrecht

## Minimax algorithm

```

function MINIMAX-DECISION(game) returns an operator
  for each op in OPERATORS[game] do
    VALUE[op] ← MINIMAX-VALUE(APPLY(op, game), game)
  end
  return the op with the highest VALUE[op]

```

```

function MINIMAX-VALUE(state, game) returns a utility value
  if TERMINAL-TEST[game](state) then
    return UTILITY[game](state)
  else if MAX is to move in state then
    return the highest MINIMAX-VALUE of SUCCESSORS(state)
  else
    return the lowest MINIMAX-VALUE of SUCCESSORS(state)


```

CS 1571 Intro to AI

M. Hauskrecht

## Complexity of the minimax algorithm

- We need to explore the complete game tree before making the decision




CS 1571 Intro to AI

M. Hauskrecht

## Complexity of the minimax algorithm

- We need to explore the complete game tree before making the decision



- Impossible for large games
  - Chess: 35 operators, game can have 50 or more moves

CS 1571 Intro to AI

M. Hauskrecht



## Solution to the complexity problem

**Two solutions:**

**1. Dynamic pruning of redundant branches** of the search tree

- identify a provably suboptimal branch of the search tree before it is fully explored
- Eliminate the suboptimal branch

**Procedure: Alpha-Beta pruning**


**2. Early cutoff of the search tree**

- uses imperfect minimax value estimate of non-terminal states (positions)

## Alpha beta pruning

- Some branches will never be played by rational players since they include sub-optimal decisions (for either player)


## Alpha beta pruning. Example



CS 1571 Intro to AI

M. Hauskrecht


## Alpha beta pruning. Example



CS 1571 Intro to AI

M. Hauskrecht


### Alpha beta pruning. Example



CS 1571 Intro to AI

M. Hauskrecht


### Alpha beta pruning. Example



CS 1571 Intro to AI

M. Hauskrecht

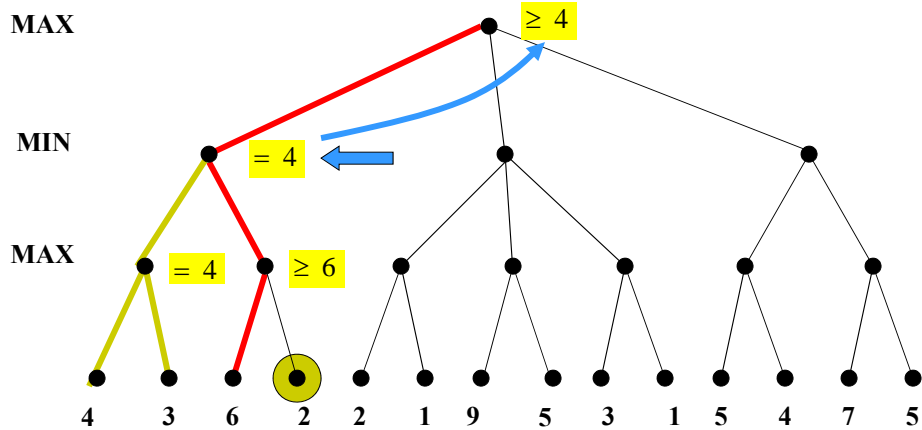
## Alpha beta pruning. Example



CS 1571 Intro to AI

M. Hauskrecht


## Alpha beta pruning. Example



CS 1571 Intro to AI

M. Hauskrecht


## Alpha beta pruning. Example



CS 1571 Intro to AI

M. Hauskrecht


## Alpha beta pruning. Example



CS 1571 Intro to AI

M. Hauskrecht


## Alpha beta pruning. Example



CS 1571 Intro to AI

M. Hauskrecht

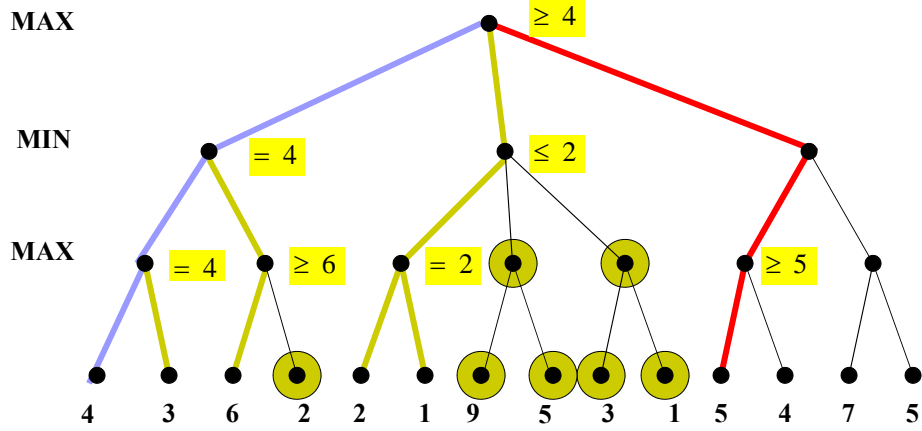
## Alpha beta pruning. Example



CS 1571 Intro to AI

M. Hauskrecht


## Alpha beta pruning. Example



CS 1571 Intro to AI

M. Hauskrecht

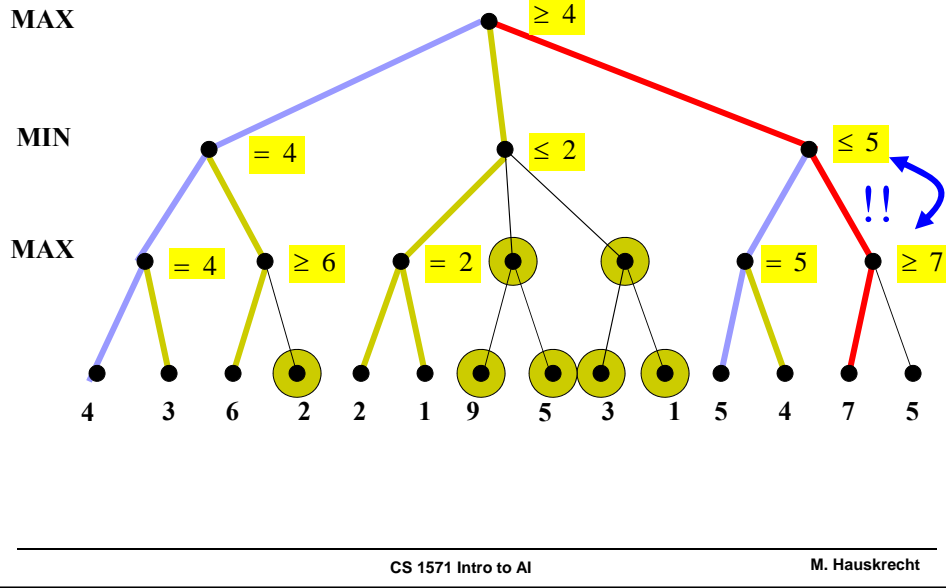
## Alpha beta pruning. Example



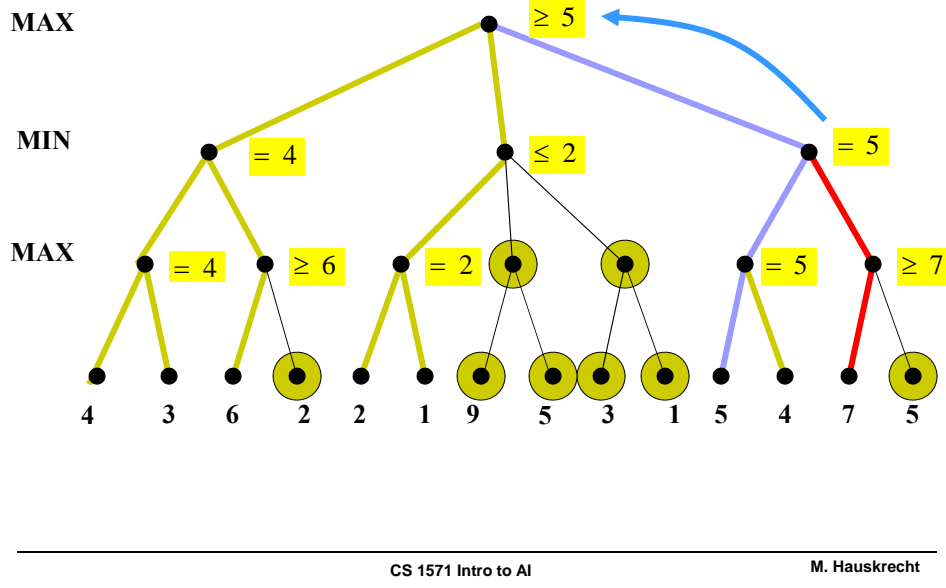
CS 1571 Intro to AI

M. Hauskrecht

### Alpha beta pruning. Example

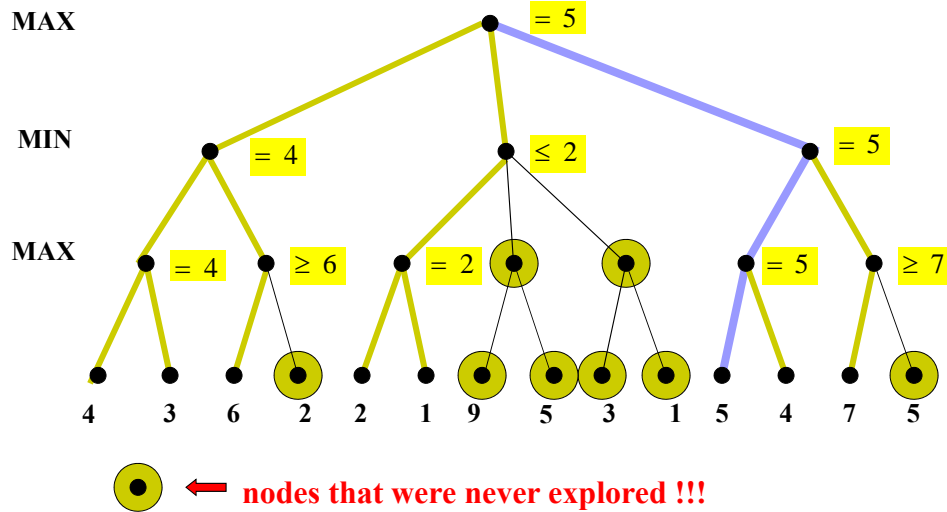


### Alpha beta pruning. Example





## Alpha beta pruning. Example



CS 1571 Intro to AI

M. Hauskrecht

## Alpha-Beta pruning

**function** MAX-VALUE(*state*, *game*,  $\alpha$ ,  $\beta$ ) **returns** the minimax value of *state*

**inputs:** *state*, current state in game  
*game*, game description  
 $\alpha$ , the best score for MAX along the path to *state*  
 $\beta$ , the best score for MIN along the path to *state*

```

if GOAL-TEST(state) then return EVAL(state)
for each s in SUCCESSORS(state) do
   $\alpha \leftarrow$  MAX( $\alpha$ , MIN-VALUE(s, game,  $\alpha$ ,  $\beta$ ))
  if  $\alpha \geq \beta$  then return  $\beta$ 
end
return  $\alpha$ 
  
```

**function** MIN-VALUE(*state*, *game*,  $\alpha$ ,  $\beta$ ) **returns** the minimax value of *state*

```


if GOAL-TEST(state) then return EVAL(state)
for each s in SUCCESSORS(state) do
   $\beta \leftarrow$  MIN( $\beta$ , MAX-VALUE(s, game,  $\alpha$ ,  $\beta$ ))
  if  $\beta \leq \alpha$  then return  $\alpha$ 
end
return  $\beta$ 
  
```

CS 1571 Intro to AI

M. Hauskrecht

## Using minimax value estimates

- **Idea:**
  - Cutoff the search tree before the terminal state is reached
  - Use imperfect estimate of the minimax value at the leaves
    - Heuristic evaluation function



CS 1571 Intro to AI

M. Hauskrecht

## Design of evaluation functions

- **Heuristic estimate** of the value for a sub-tree
- **Examples of a heuristic functions:**
  - **Material advantage in chess, checkers**
    - Gives a value to every piece on the board, its position and combines them
  - More general **feature-based evaluation function**
    - Typically a linear evaluation function:

$$f(s) = f_1(s)w_1 + f_2(s)w_2 + \dots + f_k(s)w_k$$

$f_i(s)$  - a feature of a state  $s$

$w_i$  - feature weight

CS 1571 Intro to AI

M. Hauskrecht

## Further extensions to real games

- Restricted set of moves to be considered under **the cutoff level** to reduce branching and improve the evaluation function
  - E.g., consider only the capture moves in chess

