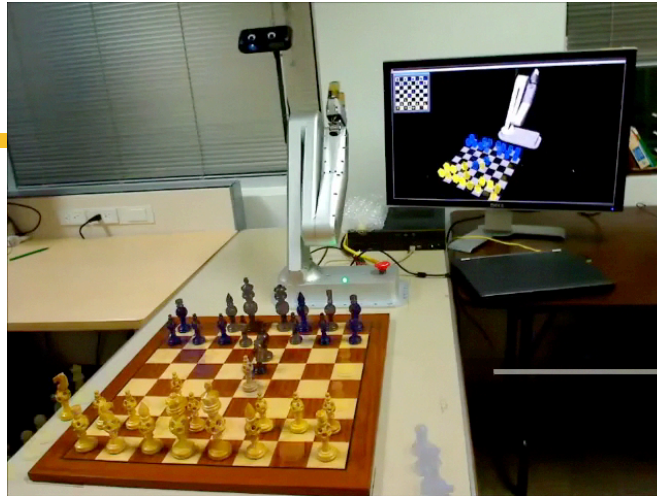


Game Playing

(aka Adversarial Search)

Ch. 5.1-5.3, 5.4.1, 5.5



Based on slides by Marie desJardin, Francisco Iacobelli, Dieter Fox

1

Bookkeeping

- HW 2 out; please look it over (shorter than HW1)
- We won't do further CSPs in class but make sure to do the reading
- Today: Game playing/search in multi-player games
 - Framework
 - Minimax
 - Alpha-beta pruning
 - Expectiminimax (games with chance elements)

2

Why Study Games?

- Clear criteria for success
- Offer an opportunity to study problems involving {hostile / adversarial / competing} agents.
- Interesting, hard problems which require minimal setup
- Often define very large search spaces
 - Chess has 35^{100} nodes in search tree, 10^{40} legal states
- Many problems can be formalized as games

3

State-of-the-art

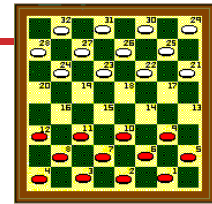
- Checkers: “Chinook”, an AI program with a very large endgame database, is world champion and can provably never be beaten. Retired 1995.
- Chess:
 - Deep Blue beat Gary Kasparov in 1997
 - Garry Kasparov vs. Deep Junior (Feb 2003): tie!
 - Kasparov vs. X3D Fritz (November 2003): tie!
 - Deep Fritz beat world champion Vladimir Kramnik (2006)
 - Now computers play computers

4

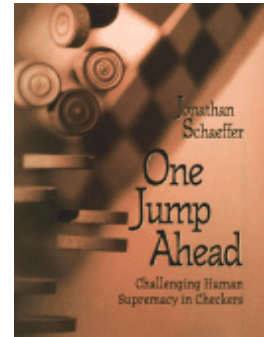
Chinook

- World Man-Machine Checkers Champion, developed by researchers at the University of Alberta.
- Earned this title by competing in human tournaments, winning the right to play for the world championship, eventually defeating the best players in the world.
- Play it! <http://www.cs.ualberta.ca/~chinook>
- Developers have fully analyzed the game of checkers, and can provably never be beaten
 - <http://www.sciencemag.org/cgi/content/abstract/1144079v1>

The board set for play



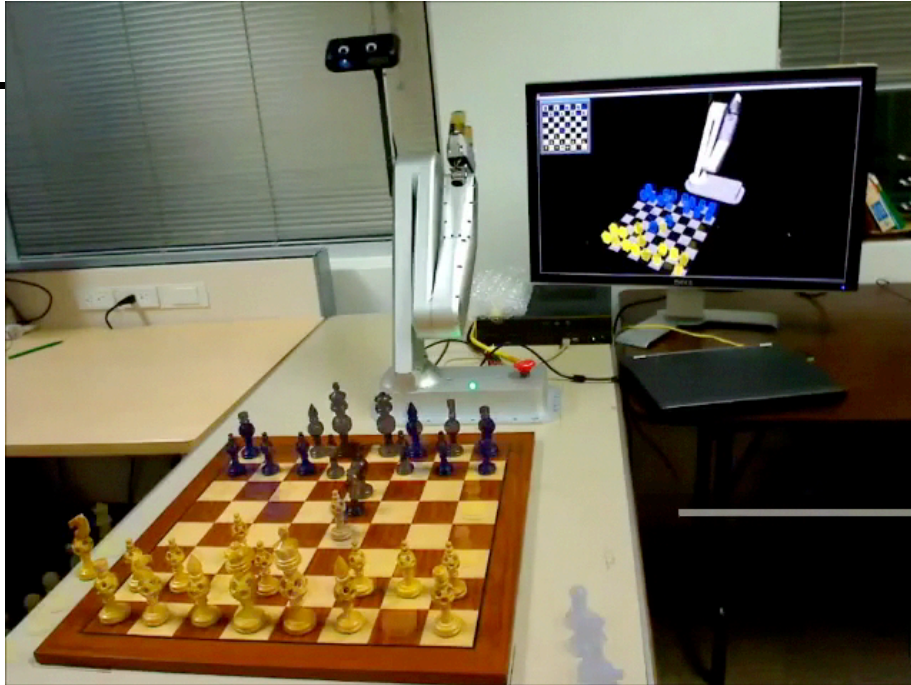
Red to play



5



6



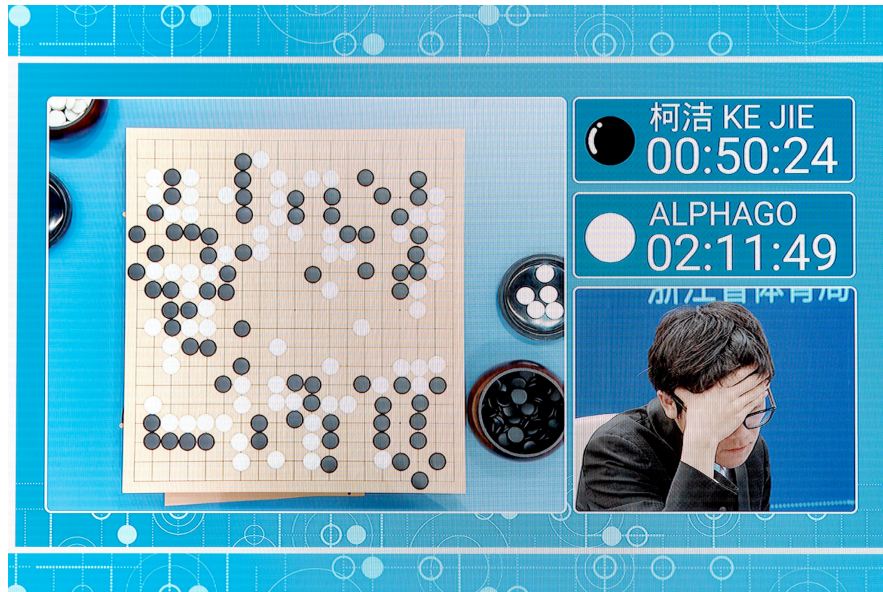
7

State-of-the-art: Go

- Computers finally got there: **AlphaGo!**
 - Made by Google DeepMind in London
- 2015: Beat a professional Go player without handicaps
- 2016: Beat a 9-dan professional without handicaps
- 2017: Beat Ke Jie, #1 human player
- 2017: DeepMind published AlphaGo Zero
 - No human games data
 - Learns from playing itself
 - Better than AlphaGo in 3 days of playing

8

AlphaGo Master defeated Ke Jie by three to zero during its 60 straight wins in the online games at the end of 2016 and beginning of 2017.



www.wired.com/2017/05/noggles-alpha-go-levels-board-games-power-grid

9

State-of-the-art: Bridge

- Bridge: “Expert-level” AI, but no world champions... exactly
 - Bridge is stochastic: the computer has imperfect information.
 - 2006: “computer bridge world champion Jack played seven top Dutch pairs ... and two reigning European champions. A total of 196 boards were played. ... Overall, the program lost by a small margin (359 versus 385).”
- 2022: NukKAI’s bridge-playing computer Nook defeats eight world bridge champions in Paris (playing, but not bidding)

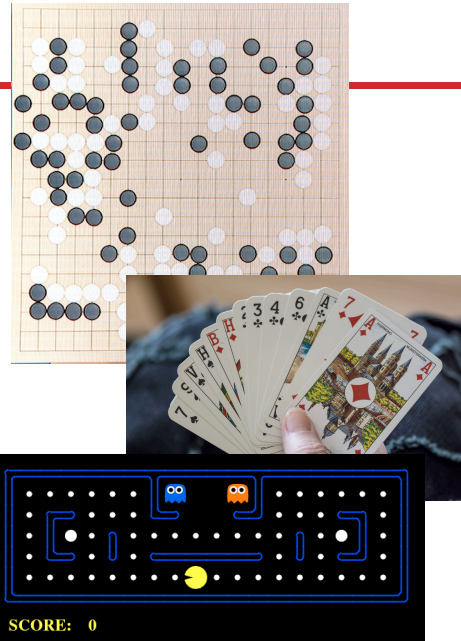


pxhere.com/en/photo/815217
wikipedia: Computer_bridge

10

Game Playing

- Many different kinds of games!
- Choices:
 - Deterministic or stochastic?
 - One, two, or more players?
 - Perfect information (can you see the state)?
- We want algorithms for calculating a **strategy (policy)** that recommends a move in each state



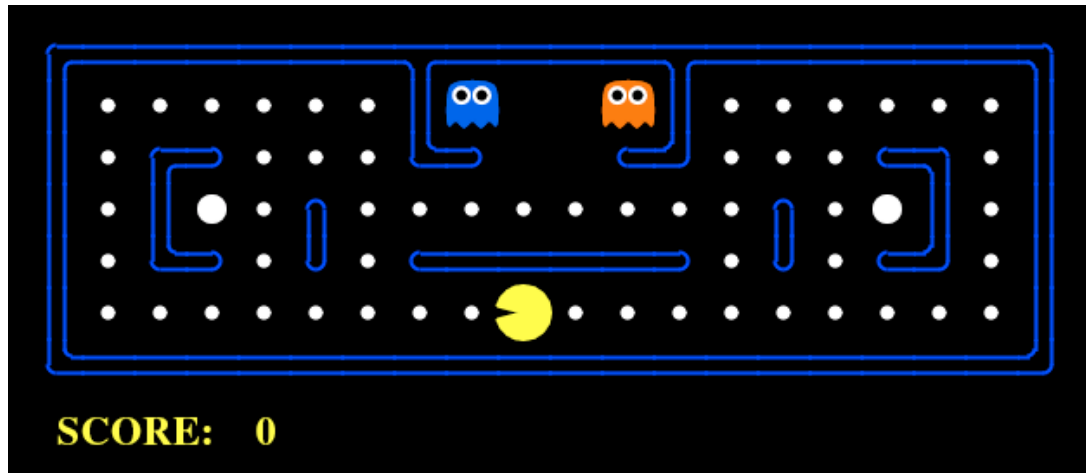
12

Typical Games

- 2-person game
- Players alternate moves
- Easiest games are:
 - **Zero-sum:** one player's loss is the other's gain
 - **Fully observable:** both players have access to complete information about the state of the game.
 - **Deterministic:** No chance (e.g., dice) involved
- Tic-Tac-Toe, Checkers, Chess, Go, Nim, Othello
- Not: Bridge, Solitaire, Backgammon, ...

13

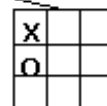
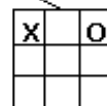
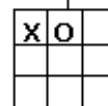
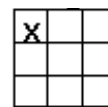
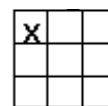
Adversarial Search



14

The Core Idea

- Searching for moves to make
- **Interleaved** with changes to the world state (from opponents' move)
- I make this move:
- Then, one of the following happens
- We don't control which, so we need to consider those possibilities when figuring out the next move after that

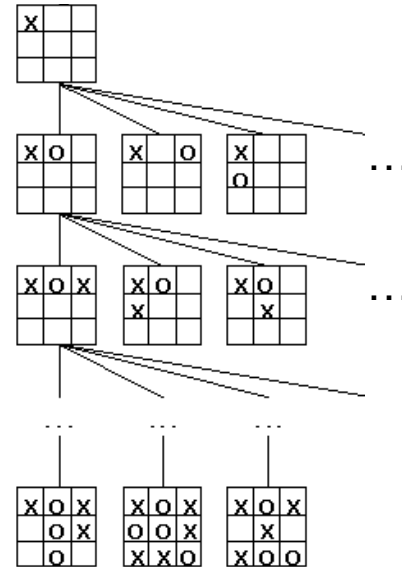


...

15

The Core Idea

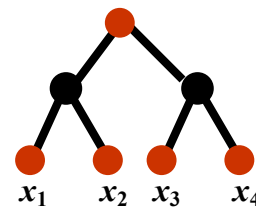
- We need to consider forward-looking possibilities when planning the next move
 - If I make this move, they might move here, here, or here; then I could ...
- Or do we?
 - We can't hold the entire search space in memory
 - So what to do?



16

How to Play (How to Search)

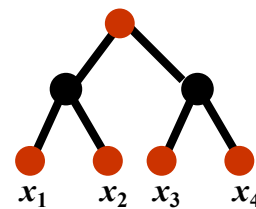
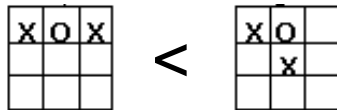
- Naïve approach:
 - From current game state:
 1. Consider all the legal moves you can make
 2. Compute new position resulting from each move
 3. Evaluate each resulting position
 4. Decide which is best
 5. Make that move
 6. Wait for your opponent to move
 7. Repeat
- Evaluating states rather than looking ahead



17

How to Play (How to Search)

- Key problems:
 - Representing the “board” (game state)
 - We’ve seen that there are different ways to make these choices
 - Generating all legal next boards
 - That can get ugly
- **Evaluating a position**



18

Evaluation Function

- **Evaluation function** or **static evaluator** is used to evaluate the “goodness” of a game *position* (state)
- Zero-sum assumption allows *one* evaluation function to describe goodness of a board for *both* players
 - One player’s gain of n means the other loses n
 - How?

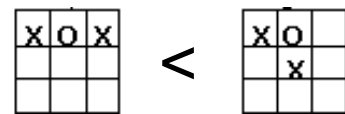


Photograph: Thanakorn Suppamethasawat/EyeEm/Getty Images

19

Evaluation Function: The Idea

- I am always trying to reach the **highest** value
- **You** are always trying to reach the **lowest** value
- Captures everyone's goal in a single function
- Sample evaluation function:
 - $f(n) > 0$: position n good for me and bad for you
 - $f(n) < 0$: position n bad for me and good for you
 - $f(n) = 0 \pm \epsilon$: position n is a neutral position
 - $f(n) = +\infty$: win for me
 - $f(n) = -\infty$: win for you



X wants high values (max)
O wants low values (min)

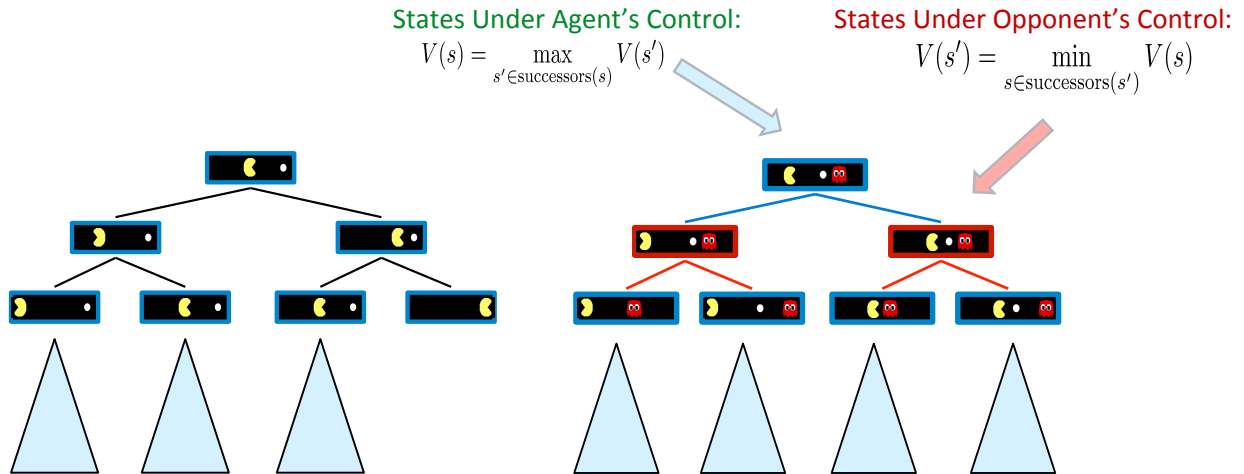
20

Evaluation Function Examples

- Example of an evaluation function for Tic-Tac-Toe:
 - $f(n) = [\text{\#3-lengths open for X}] - [\text{\#3-lengths open for O}]$
 - A 3-length is a complete row, column, or diagonal
- Alan Turing's function for chess
 - $f(n) = w(n)/b(n)$
 - $w(n)$ = sum of the point value of white's pieces
 - $b(n)$ = sum of black's
- Core idea: one player is trying to **maximize** and opponent is trying to **minimize** some evaluation function

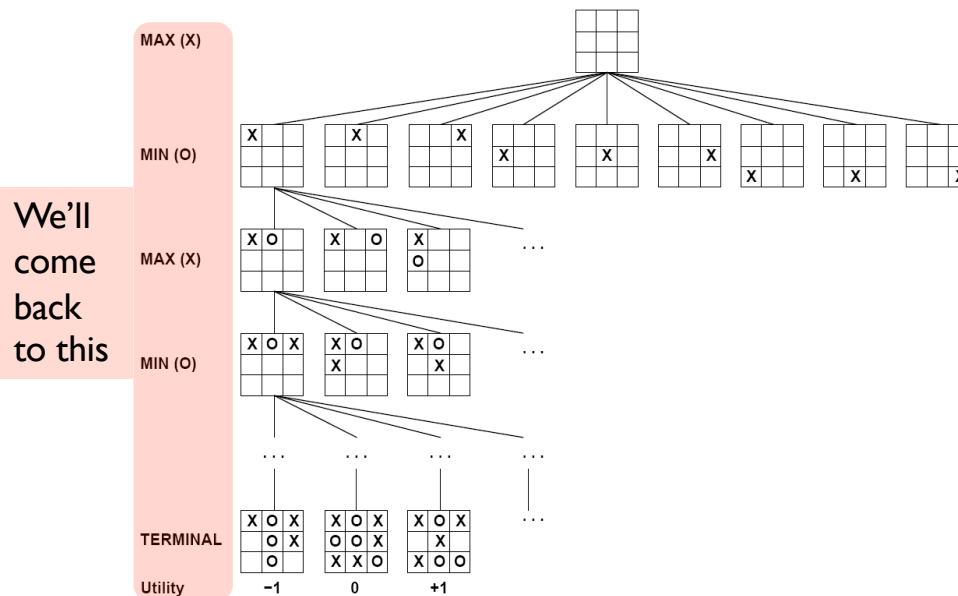
21

One- vs. Two-Player Games



22

Example Game Search Space



23

Evaluation function examples

- Most evaluation functions are specified as a **weighted sum** of position features:

$$f(n) = w_1 * feat_1(n) + w_2 * feat_2(n) + \dots + w_n * feat_k(n)$$

- Example features for chess: piece count, piece placement, squares controlled, ...
- Deep Blue had over **8000** features in its nonlinear evaluation function!

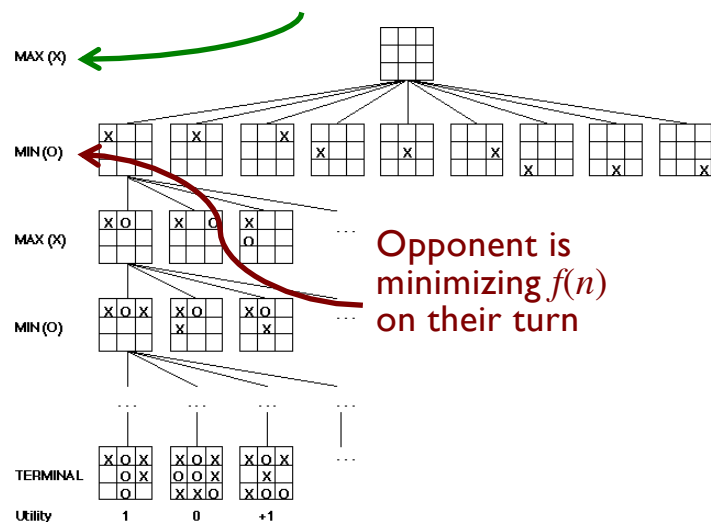
square control, rook-in-file, x-rays, king safety, pawn structure, passed pawns, ray control, outposts, pawn majority, rook on the 7th blockade, restraint, trapped pieces, color complex, ...

24

Game trees

- Problem spaces for typical games are represented as trees
- Player must decide best single move to make next
- Root node = current board configuration
- Arcs = possible legal moves for a player

I am maximizing $f(n)$ on my turn



25

Game trees

- **Static evaluator function**

- Rates a board position
- $f(\text{board}) = R$, with $f > 0$ for me, $f < 0$ for you

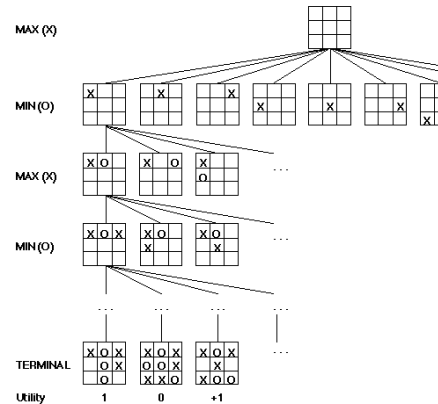
- If it is **my turn** to move:

- Root is labeled “**MAX**” node

- Otherwise it is a “**MIN**” node (**opponent's turn**)

- Each level's nodes are all MAX or all MIN

- Nodes at level i are opposite those at level $i + 1$



26

Minimax: The Intuition

- But do we have to consider **any** move an opponent might make?
- Consider the best move to make, **based on what your opponent will actually do if you make that move**
 - First apply a max function, then a min function, then a max function...
- “Look ahead”: consider the resulting board state after you make your move, and after the opponent makes their next **sensible** move
- Can consider arbitrarily far forward

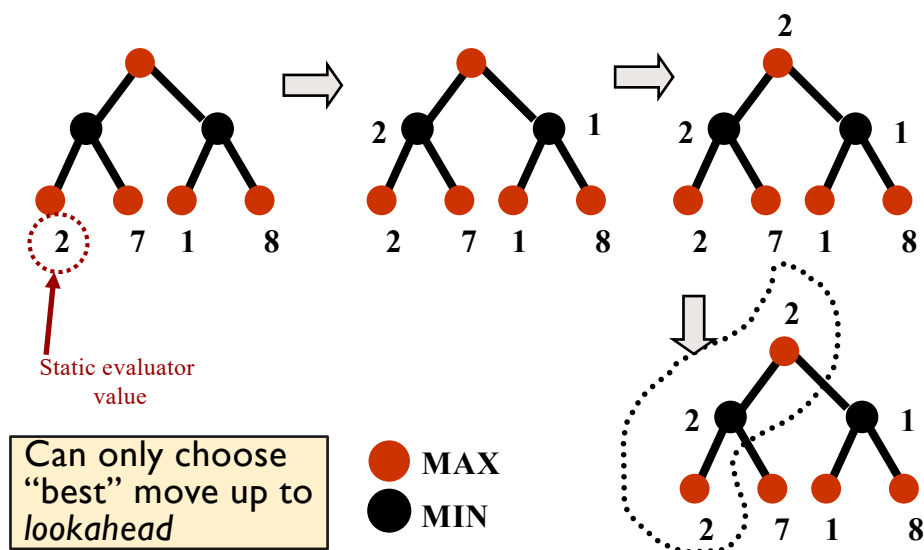
27

Minimax Procedure

- Create start node: MAX node, current board state
- Expand nodes down to a **depth** of *lookahead*
- Apply evaluation function at each leaf node
- “Back up” values for each non-leaf node until a value is computed for the root node
 - MIN: backed-up value is **lowest** of children’s values
 - MAX: backed-up value is **highest** of children’s values
- Pick operator associated with the child node whose backed-up value set the value at the root

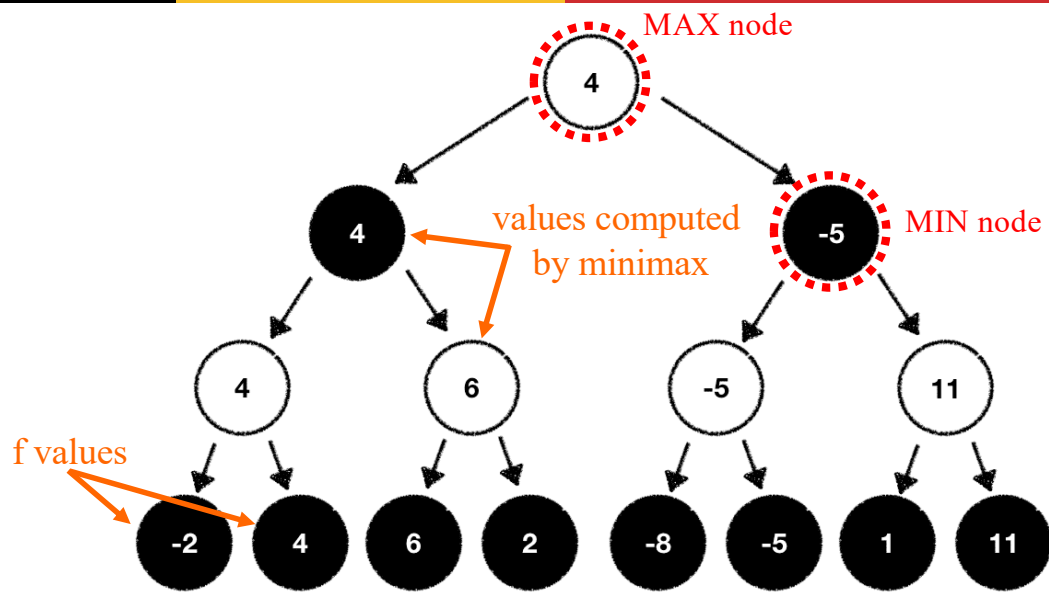
28

Minimax Algorithm



29

Example Minimax Tree

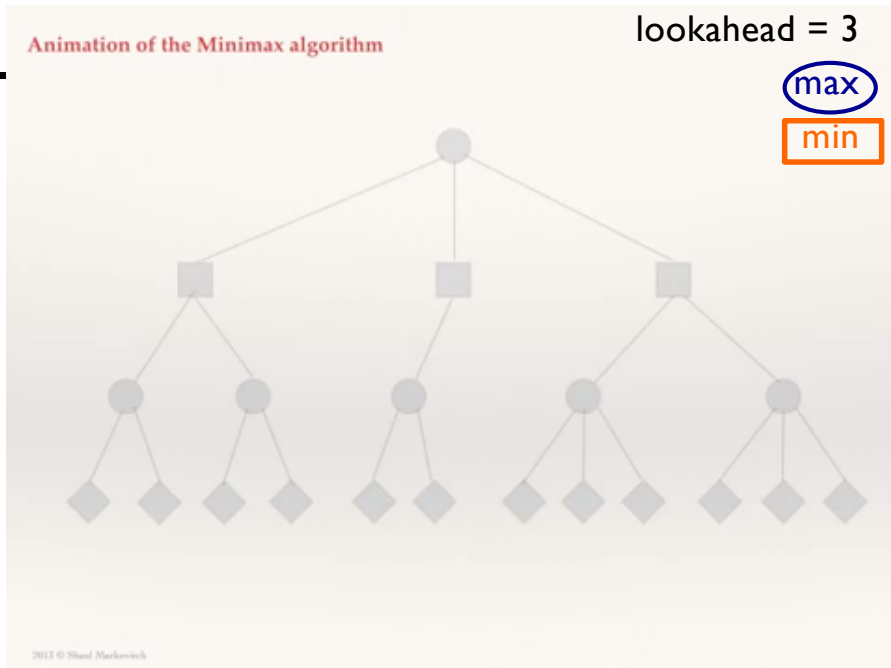


30

Animation of the Minimax algorithm

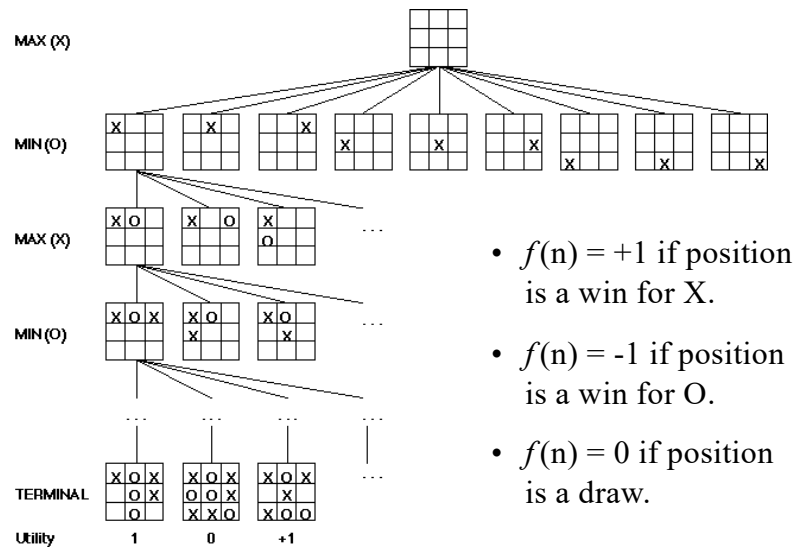
lookahead = 3

max
min



31

Partial Game Tree for Tic-Tac-Toe



32

Example: Nim

- In Nim, there are a certain number of objects (coins, sticks, etc.) on the table – we'll play 7-coin Nim
- Each player in turn has to pick up either one or two objects
- Whoever picks up the last object loses



34

Nim Game Tree

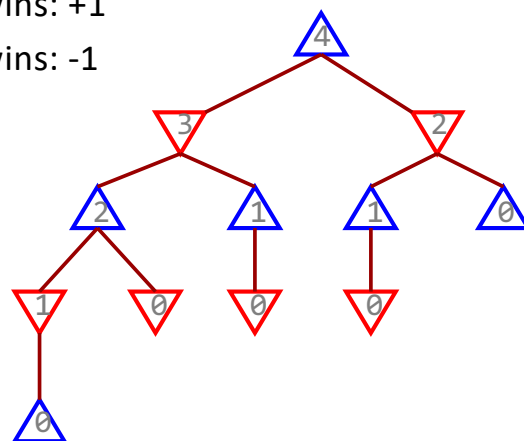
- **In-class exercise:**
- Draw minimax search tree for 4-coin Nim
- Things to consider:
 - What's your start state?
 - What's the maximum depth of the tree? Minimum?
- Pick up either one or two objects
- Whoever picks up the last object loses



35

Nim Game Tree

Player 1 wins: +1
Player 2 wins: -1

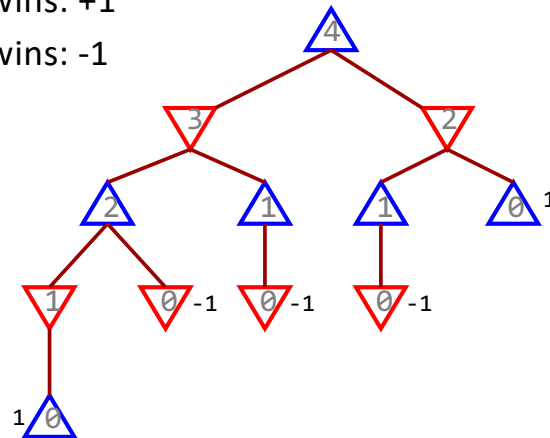


37

Nim Game Tree

Player 1 wins: +1

Player 2 wins: -1

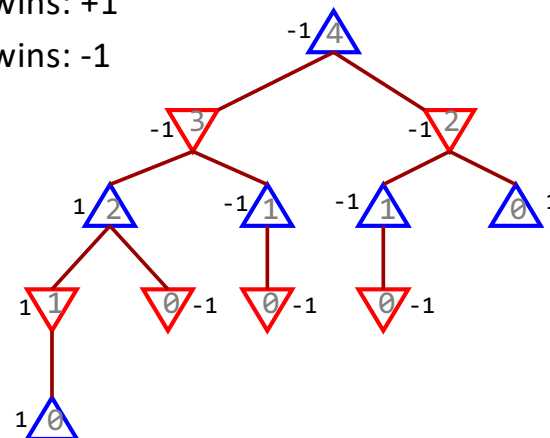


38

Nim Game Tree

Player 1 wins: +1

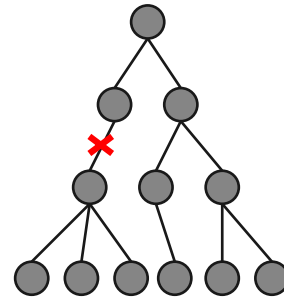
Player 2 wins: -1



39

Improving Minimax

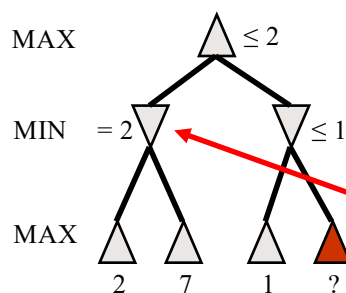
- Basic problem: must examine a number of states that is exponential in d !
- Solution: judicious **pruning** of the search tree
- “Cut off” whole sections that **can't** be part of the best solution
 - Or, sometimes, **probably won't**
 - Can be a completeness vs. efficiency tradeoff, esp. in stochastic problem spaces



40

Alpha-Beta Pruning

- We can improve on the performance of the minimax algorithm through **alpha-beta pruning**
 - Basic idea: “If you have an idea that is surely bad, don't take the time to see how truly awful it is.” – Pat Winston



- We don't need to compute the value at this node.
- No matter what it is, it can't affect the value of the root node.
- Because the MAX player will choose this value.

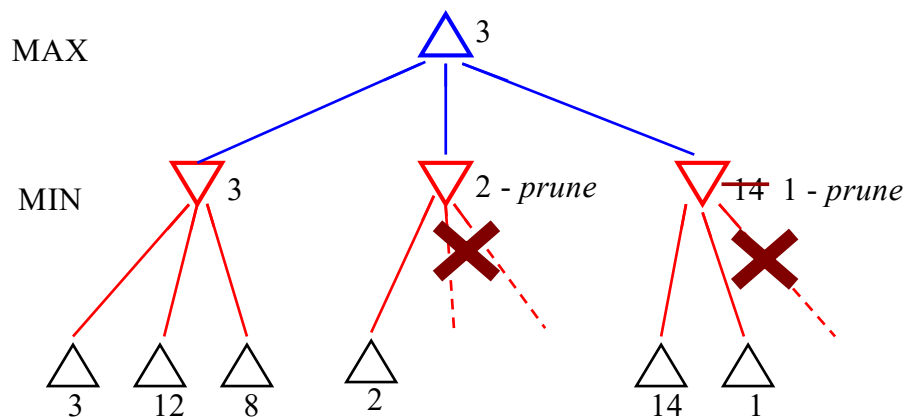
41

Alpha-Beta Pruning

- Traverse search tree in *depth-first order*
- At each **MAX** node n , $\alpha(n)$ = **maximum** value found so far
- At each **MIN** node n , $\beta(n)$ = **minimum** value found so far
 - α starts at $-\infty$ and increases, β starts at $+\infty$ and decreases
- **β -cutoff**: Given a MAX node n ,
 - Cut off search below n (i.e., don't look at any more of n 's children) if:
 - $\alpha(n) \geq \beta(i)$ for some MIN node ancestor i of n
- **α -cutoff**:
 - Stop searching below MIN node n if:
 - $\beta(n) \leq \alpha(i)$ for some MAX node ancestor i of n

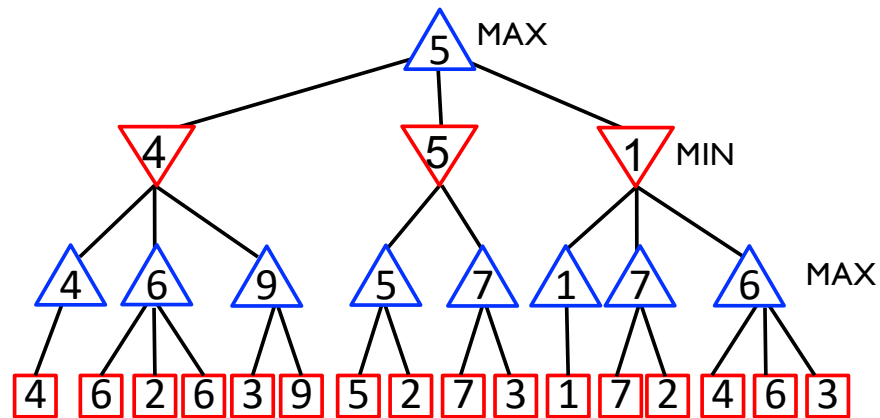
42

Alpha-Beta Example (b=3)



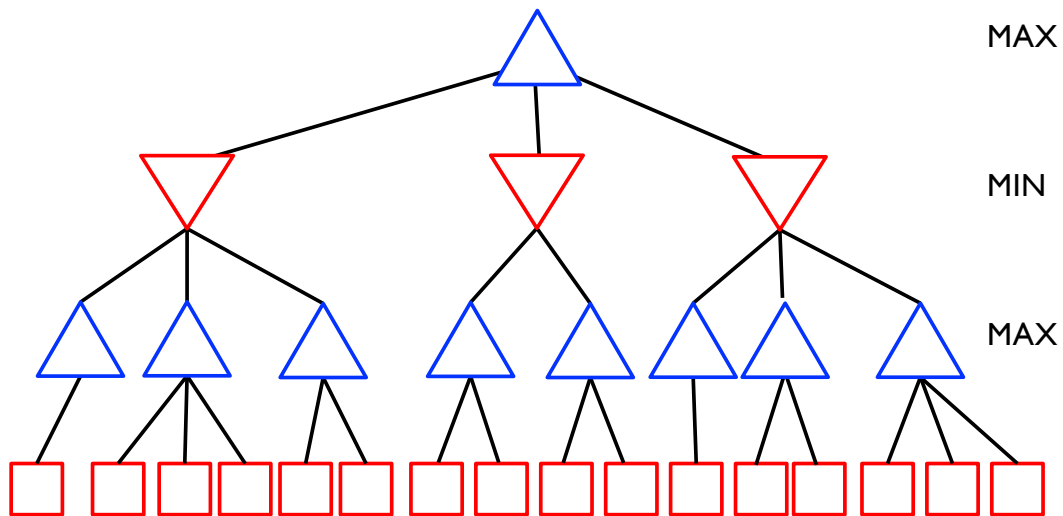
43

Alpha-Beta Pruning: Exercise



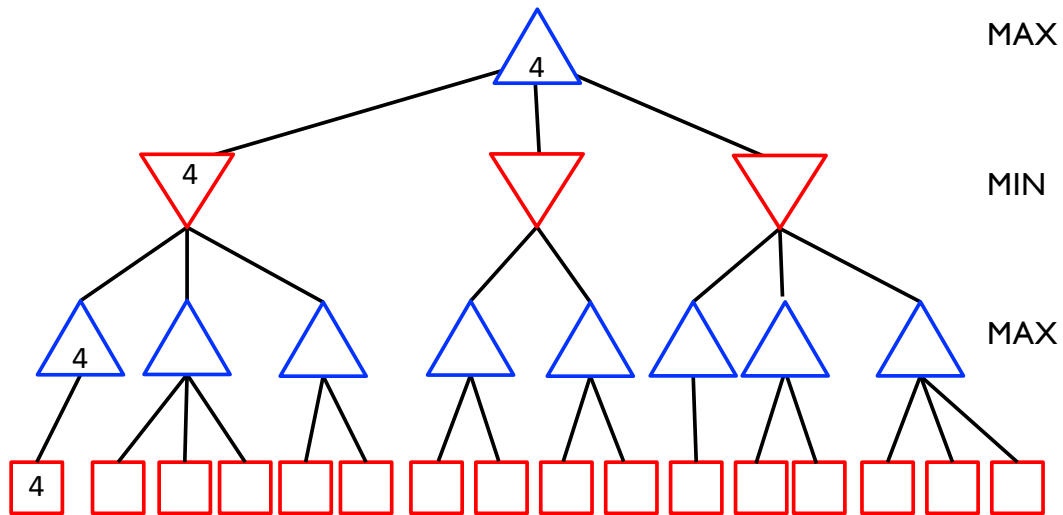
44

Alpha-beta pruning



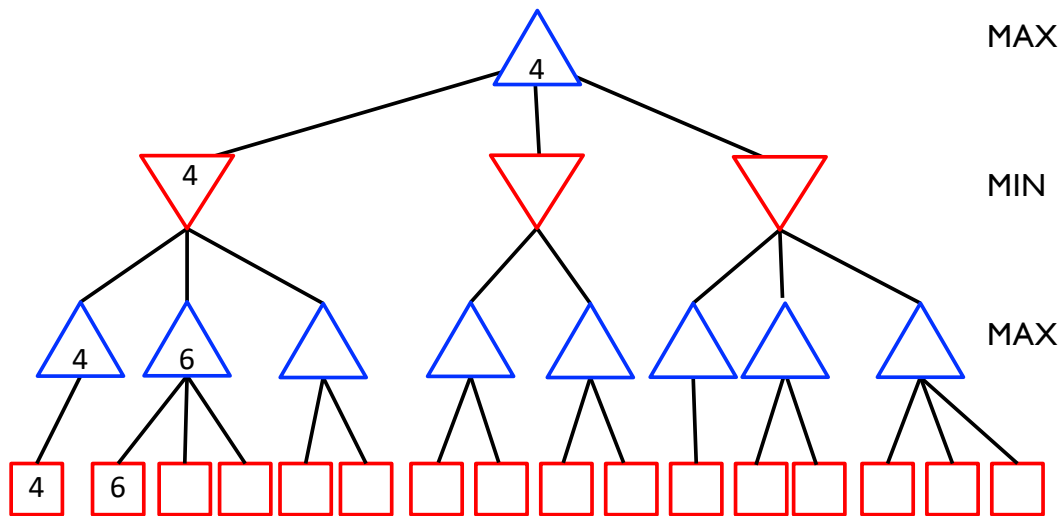
45

Alpha-beta pruning



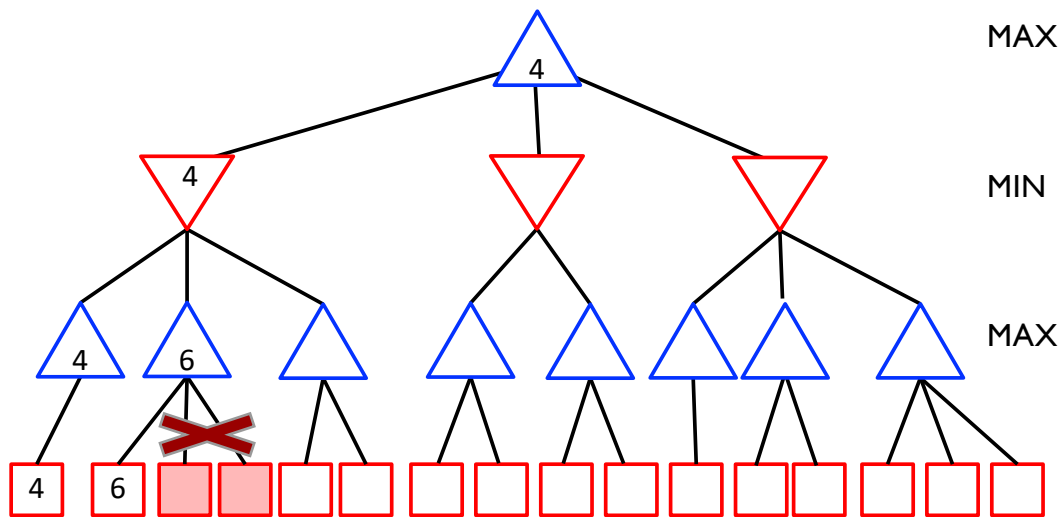
46

Alpha-beta pruning



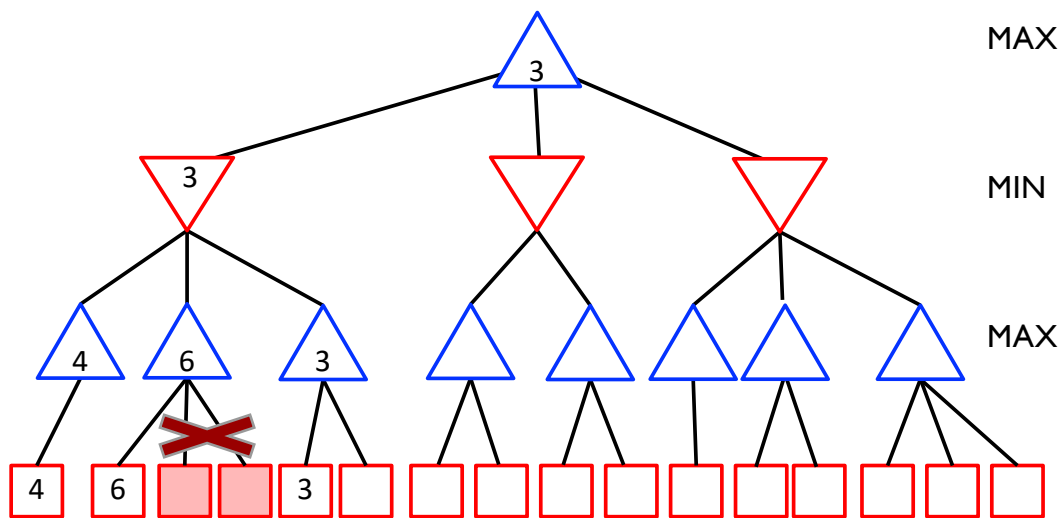
47

Alpha-beta pruning



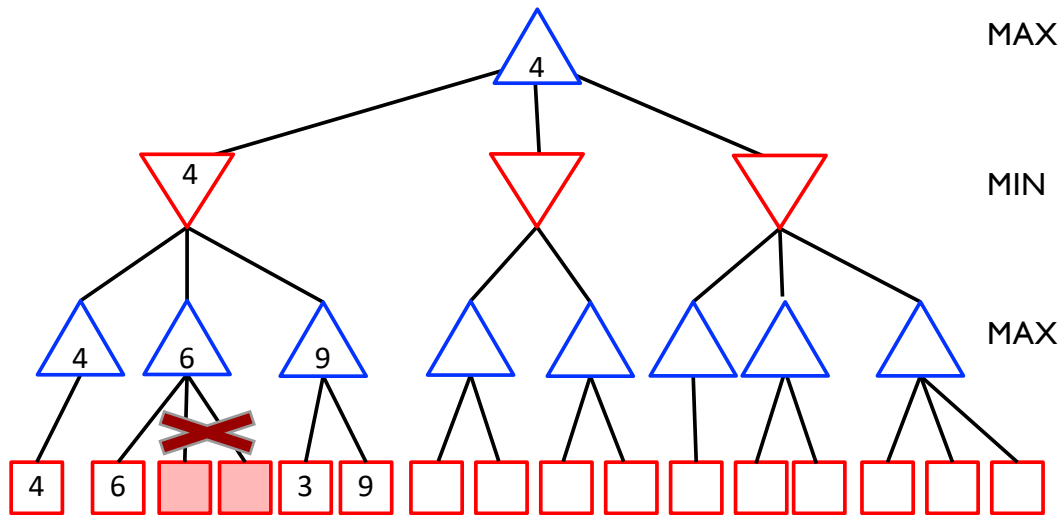
48

Alpha-beta pruning



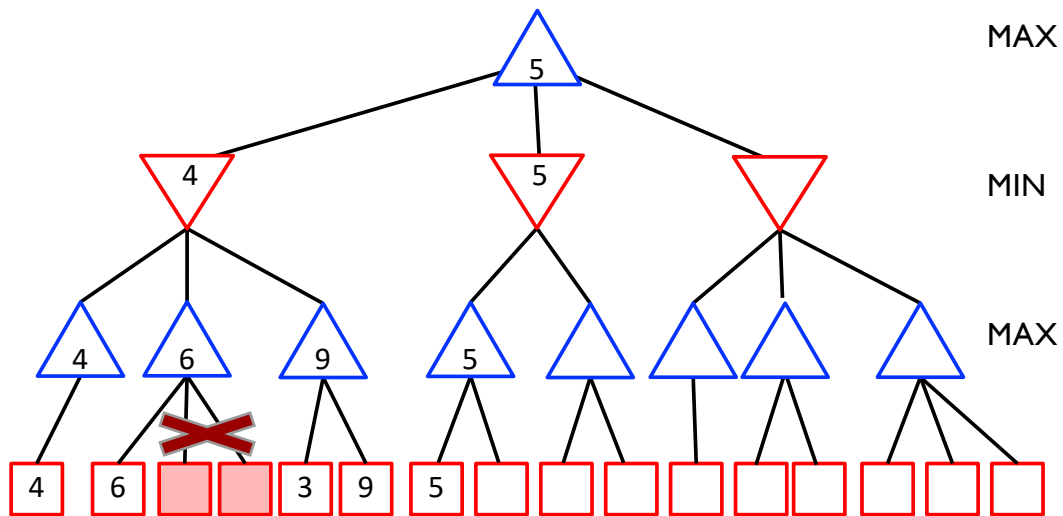
49

Alpha-beta pruning



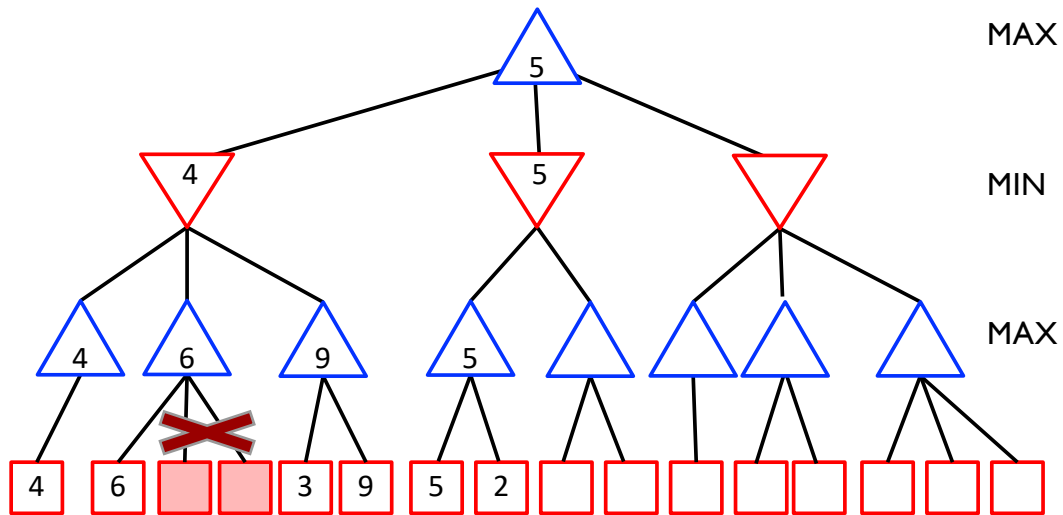
50

Alpha-beta pruning



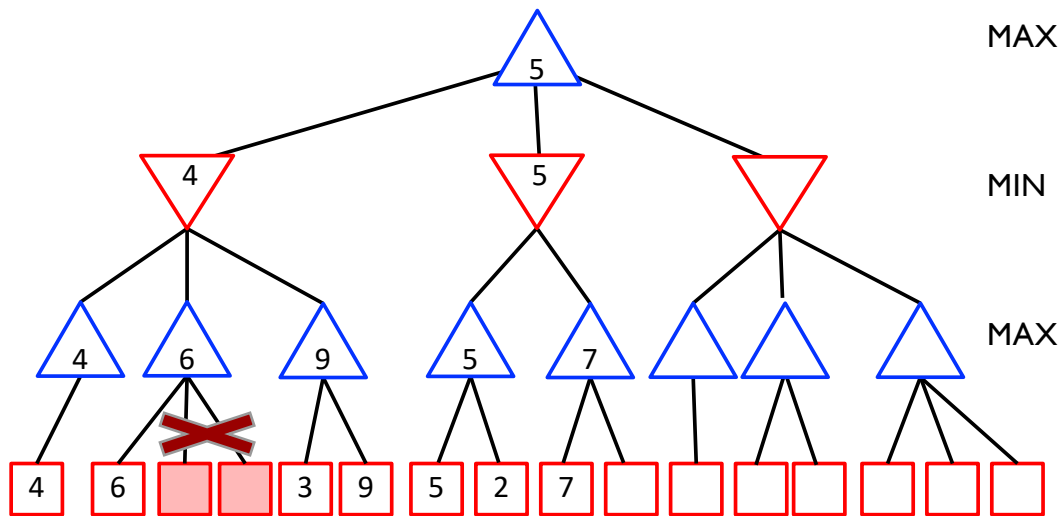
51

Alpha-beta pruning



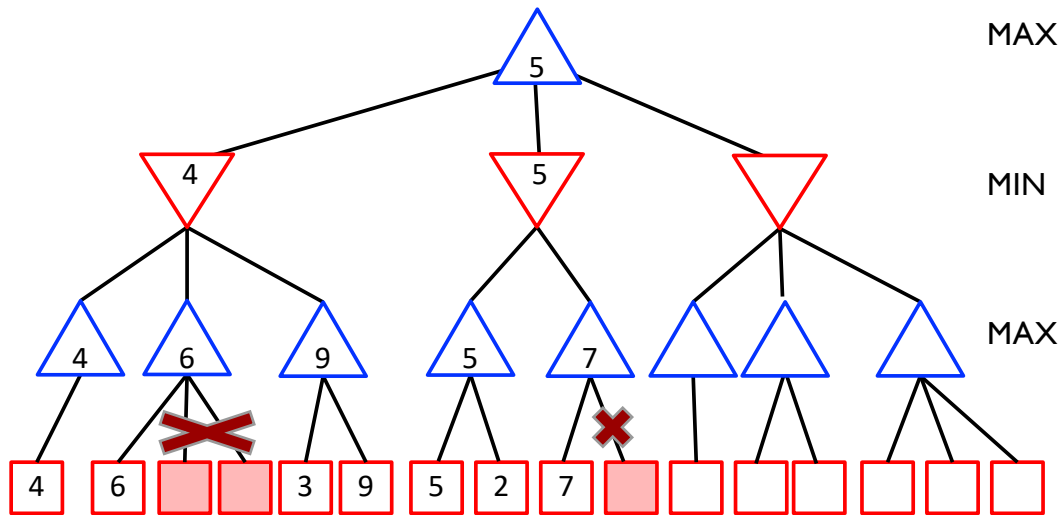
52

Alpha-beta pruning



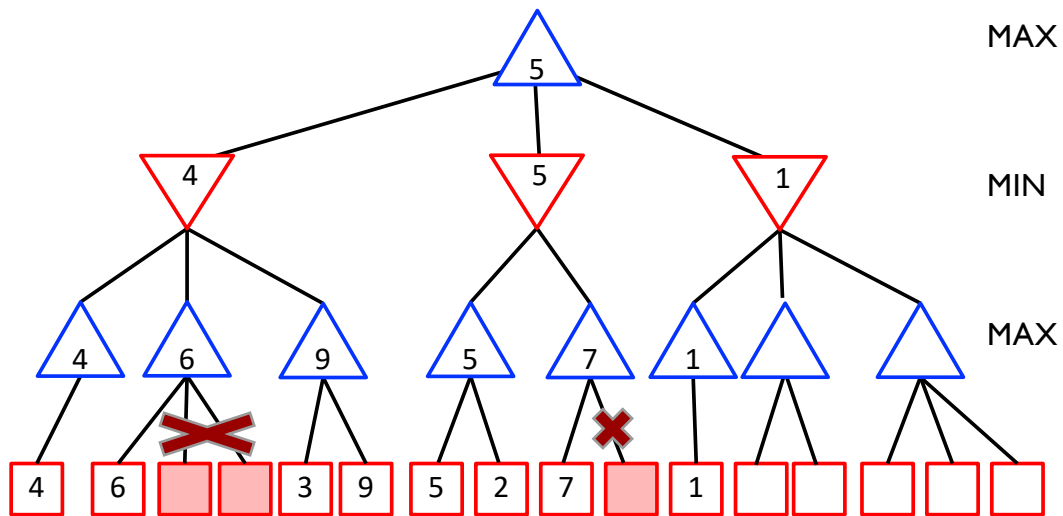
53

Alpha-beta pruning



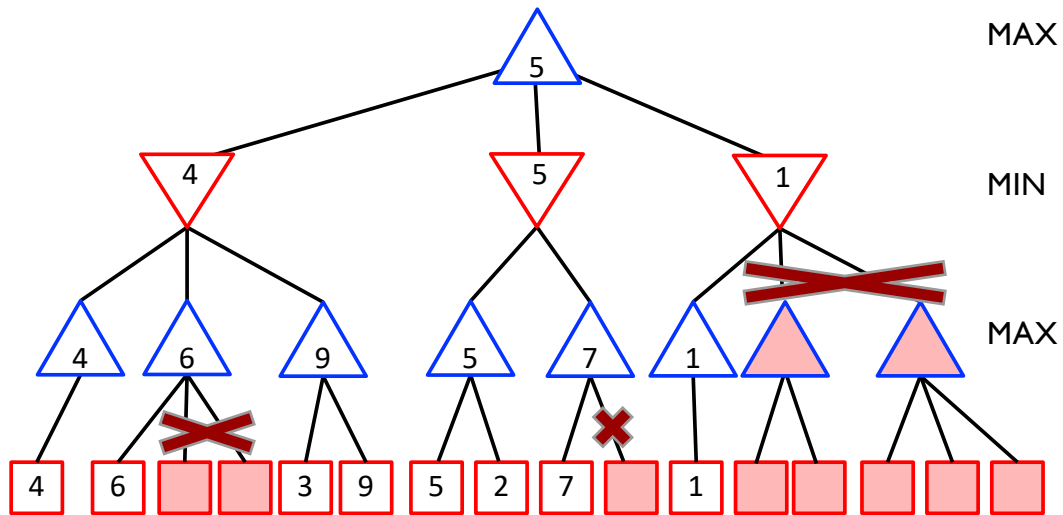
54

Alpha-beta pruning



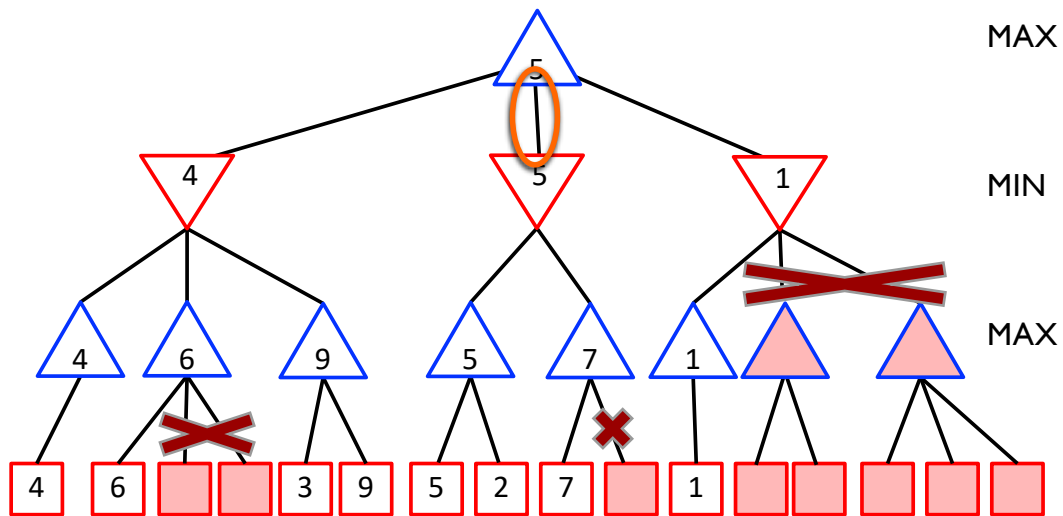
55

Alpha-beta pruning



56

Alpha-beta pruning



57

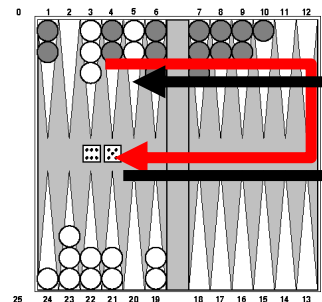
Effectiveness of Alpha-Beta

- Alpha-beta is guaranteed to:
 - Compute the same value for the root node as minimax
 - With \leq computation
- Worst case:** nothing pruned
 - Examine b^d leaf nodes
 - Each node has b children and a d -ply search is performed
- Best case:** examine only $(2b)^{d/2}$ leaf nodes.
 - So you can search twice as deep as minimax!
 - When each player's best move is the first alternative generated
- In Deep Blue, empirically, alpha-beta pruning took average branching factor from ~ 35 to ~ 6

59

Games of Chance

- Backgammon: 2-player with uncertainty
- Players roll dice to determine what moves to make
- White has just rolled 5 and 6 and has four legal moves:
 - 5-10, 5-11
 - 5-11, 19-24
 - 5-10, 10-16
 - 5-11, 11-16
- Good for decision making in adversarial problems with skill and luck

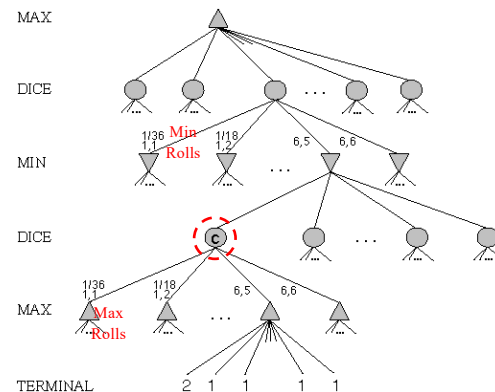


In backgammon, the dice roll determines the possible moves you can make

60

Game Trees with Chance

- **Chance nodes** (circles) represent random events
- For a random event with N outcomes:
 - Chance node has N distinct children
 - Each has a probability
 - These represent the possible states the game may be in based on the outcome of the random event
- Example:
 - Rolling 2 dice \rightarrow 21 distinct outcomes
 - Not all equally likely!



61

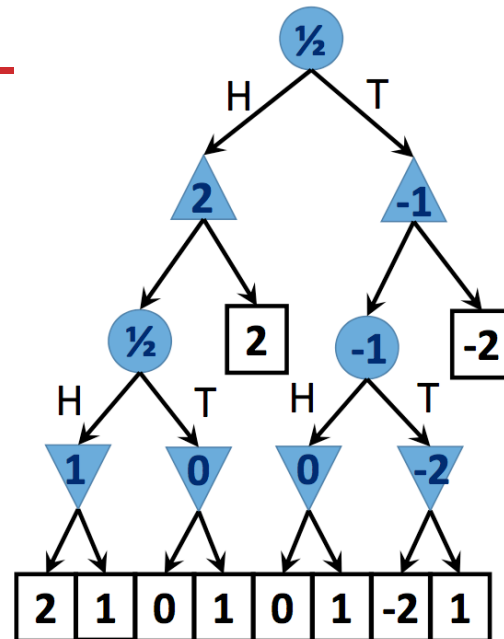
Game Trees with Chance

- For chance nodes we compute the **expected** value – the sum of the value over all possible outcomes, weighted by the likelihood of that outcome
- Use minimax to compute values for MAX and MIN nodes
- Use **expected values** for chance nodes
- $\text{Expectiminimax}(s) =$
 - $\max_a \text{Expectiminimax}(\text{Result}(s,a))$ if $\text{player}(s)=\text{max}$
 - $\min_a \text{Expectiminimax}(\text{Result}(s,a))$ if $\text{player}(s)=\text{min}$
 - $\sum_r P(r) \text{Expectiminimax}(\text{Result}(s,r))$ if $\text{player} = \text{chance}$

62

Expectiminimax Example

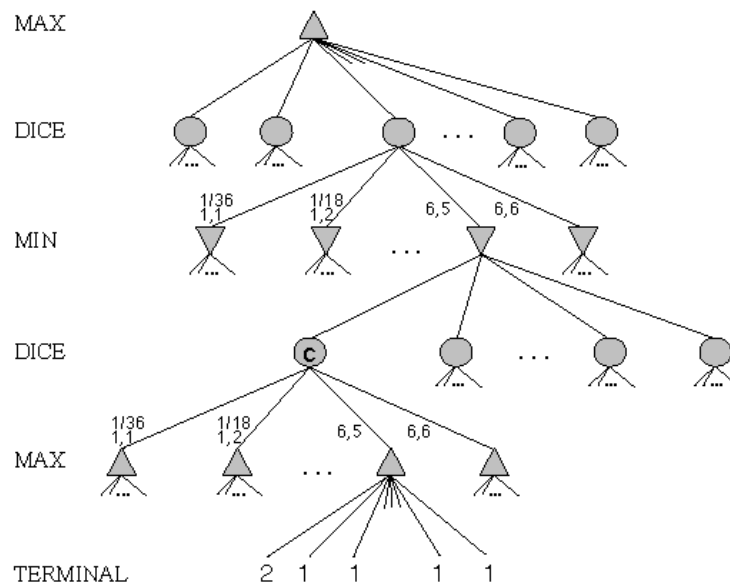
- RANDOM: Flip a coin.
- MAX: Max either stops, or continues.
 - Stop on heads: game ends, Max wins (value = \$2).
 - Stop on tails: game ends, Max loses (value = -\$2).
- RANDOM: Flip a coin.
 - HH = Min gets \$2.
 - TT = Min gets -\$2.
 - HT or TH = Min gets \$0.
- MIN: Min decides whether to keep the outcome (value as above), or pay a penalty of -\$1. Game ends.



from Lana Lazebnik

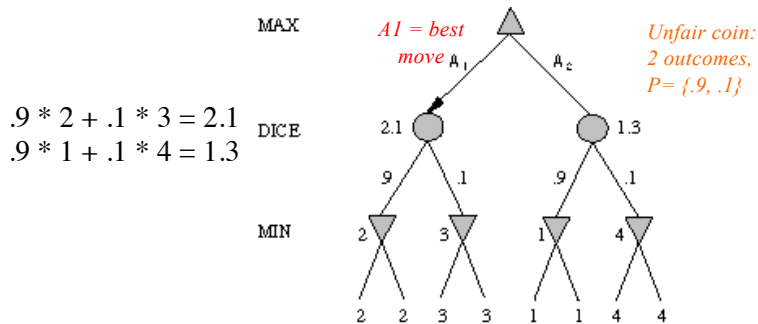
64

Game Trees with Chance



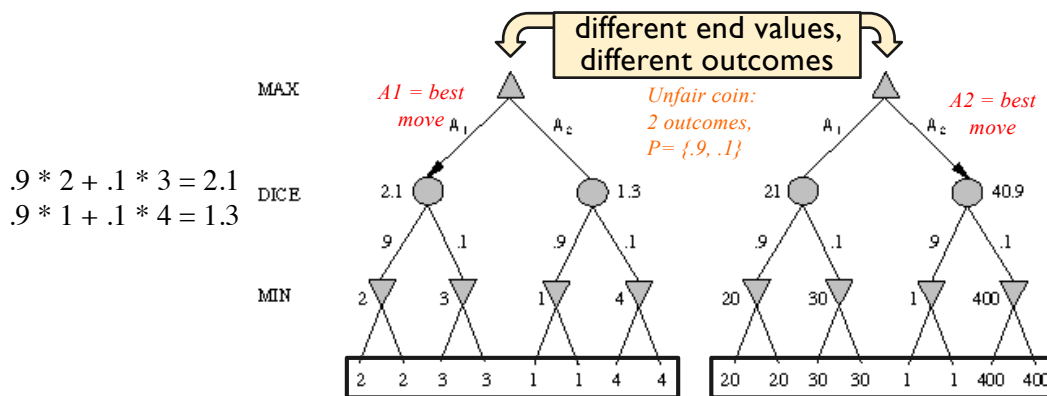
65

Meaning of the Evaluation Function



66

Meaning of the Evaluation Function



- Dealing with probabilities and expected values means being careful with “meaning” of values returned by the static evaluator
- “Relative-order preserving” (as here) change won’t change minimax, but could change the decision with chance nodes

67

Exercise: Oopsy-Nim

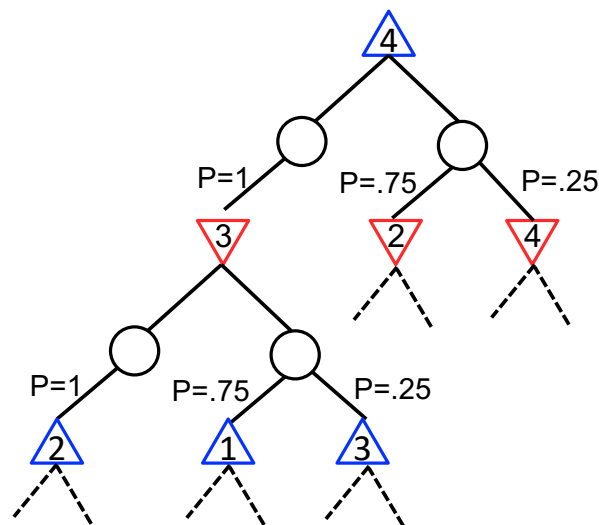
- Starts out like Nim
 - Each player in turn has to pick up either one or two objects
 - Sometimes (probability = 0.25), when you try to pick up two objects, you drop them both
 - Picking up a single object always works



- Question: Why can't we draw the entire game tree?
- Exercise: Draw the 4-ply game tree (2 moves per player)

68

Exercise: Oopsy-Nim



69