CS540 Introduction to Artificial Intelligence
Lecture 18

Young Wu
Based on lecture slides by Jerry Zhu and Yingyu Liang

July 24, 2019

Lion Game Example, Part I

Quiz (Participation)

- There are N lions, ordered by size, $i = 1, 2, 3, ..., N$, and a bunny. N takes an integer between 1 and 10 with equal probability (known to all the lions). Each lion i can choose to jump out and eat the slightly smaller lion $i - 1$, or stay hidden, and only lion 1 can eat the bunny. Each lion prefers eating to staying hungry to being eaten.
- What is the probability that the bunny is eaten?
- A: 0, B: $\frac{1}{3}$, C: $\frac{1}{2}$, D: $\frac{2}{3}$, E: 1

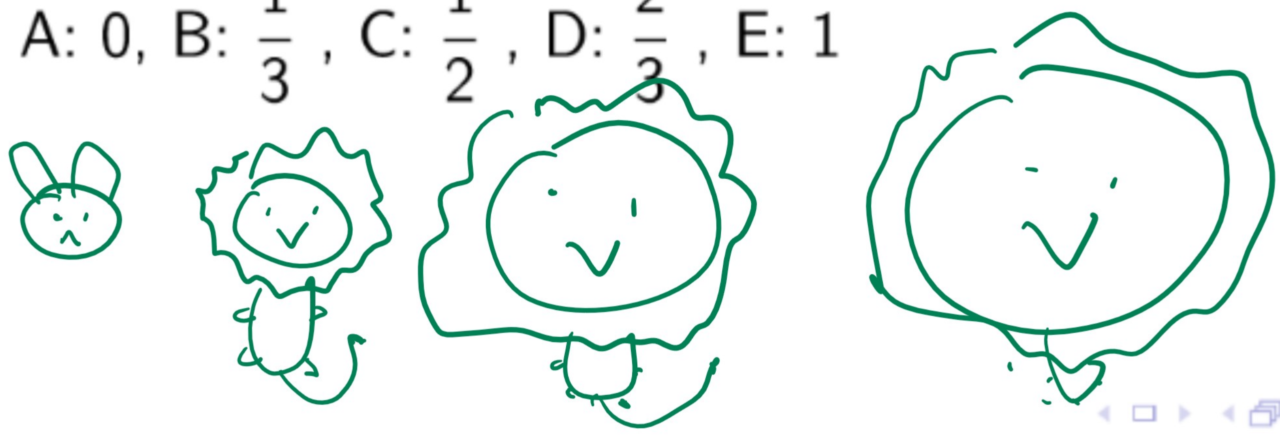

Lion Game Example, Part II

Quiz (Participation)

N = 1   B L_1 ✓

N = 2   B L_1 (hide) L_2 ✗

→ N = 3   B L_1 (eat) L_2 (hide) L_3 ✓

N = 4   B L_1 (hide) L_2 (eat) L_3 (hide) L_4

✗

odd — B eaten

even — no

$prob = \frac{1}{2}$

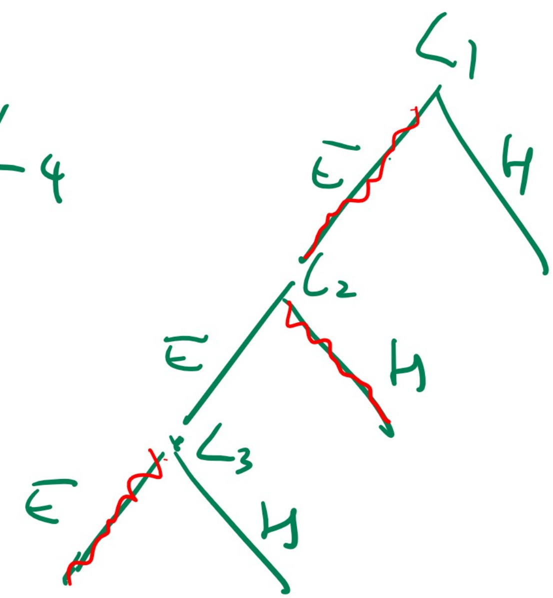

Pirate Game Example, Part I

Quiz (Participation)

- 5 pirates got 100 gold coins. Each pirate takes a turn to propose how to divide the coins, and all pirates who are still alive will vote whether to accept the proposal or reject the proposal, kill the pirate, and continue to the next round. Use strict majority rule for the vote, and use the assumption that if a pirate is indifferent, he or she will vote reject.
- How will the first pirate propose?
- A: $(0, 0, 0, 0, 100)$
- B: $(20, 20, 20, 20, 20)$
- C: $(94, 0, 1, 2, 3)$
- D: $(97, 0, 1, 0, 2)$
- E: $(98, 0, 1, 0, 1)$

Pirate Game Example, Part II

Quiz (Participation)

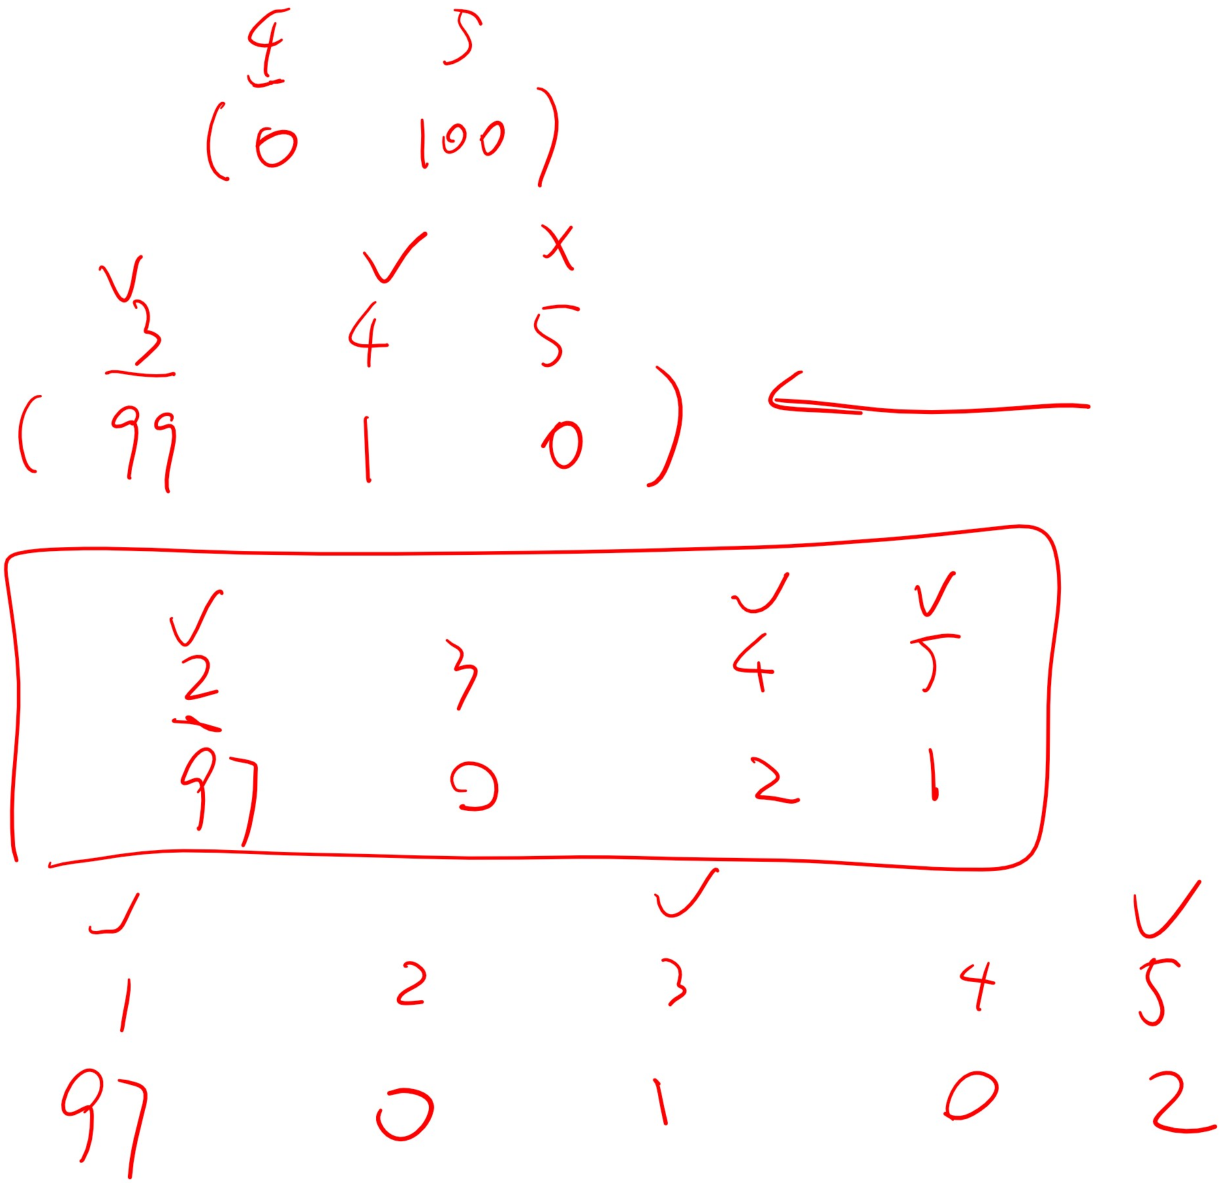

Wage Competition Example, Part I

Quiz (Participation)

- An applicant can produce 20 dollars worth of products per hour. Two profit-maximizing companies compete to hire the applicant. Company 1 makes an offer first, then company 2 sees the offer and makes another offer. The applicant goes to the company that offers a higher wage, and in case of a tie, he or she will flip a fair coin to decide.
- What should the offers be (hourly wage, only integer amounts allowed)?
- A: $(20, 20)$
- B: $(20, 19)$
- C: $(19, 20)$
- D: $(19, 19)$
- E: None of the above.

D: (19, 19)

E: none of the above.

Wage Competition Example, Part II

Quiz (Participation)

A: ~~0~~ 2 B: ~~1~~ 3

A: 19 B: 19

π = 0.5 π = 0.5

fix B, max A

fix A max B

20 − 20 = 0

(18, 18) also a solution.

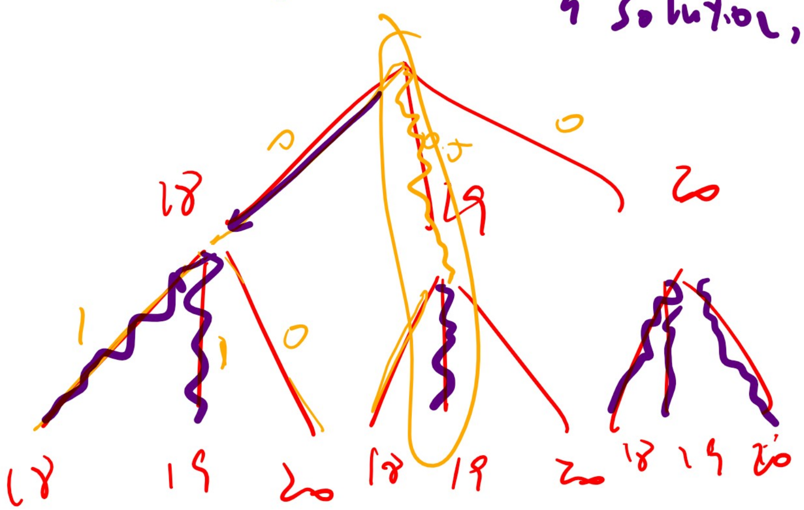

(19, 19)

Nash equilibrium → next week.

Game Tree

Motivation

- The initial state is the beginning of the game.
- There are no goal states, but there are multiple terminal states in which the game ends.
- Each successor of a state represents a feasible action (or a move) in the game.
- The search problem is to find the terminal state with the lowest cost (or usually the highest reward).

Adversarial Search

Motivation

- The main difference between finding solutions of games and standard search problems or local search problems is that part of the search is performed by an opponent adversarially.
- Usually, the opponent wants to maximize the cost or minimize the reward from the search. This type of search problems is called adversarial search.
- In game theory, the solution of a game is called an equilibrium. It is a path in which both players do not want to change actions.

next week

Backward Induction

Motivation

- Games are usually solved backward starting from the terminal states.
- Each player chooses the best action (successor) given the (already solved) optimal actions of all players in the subtrees (called subgames).

Zero-Sum Games

Motivation

- If the sum of the reward or cost over all players at each terminal state is 0, the game is called a zero-sum game.
- Usually, for games with one winner: the reward for winning and the cost of losing are both 1. If the game ends with a tie, both players get 0.

Tic Tac Toe Example

Motivation

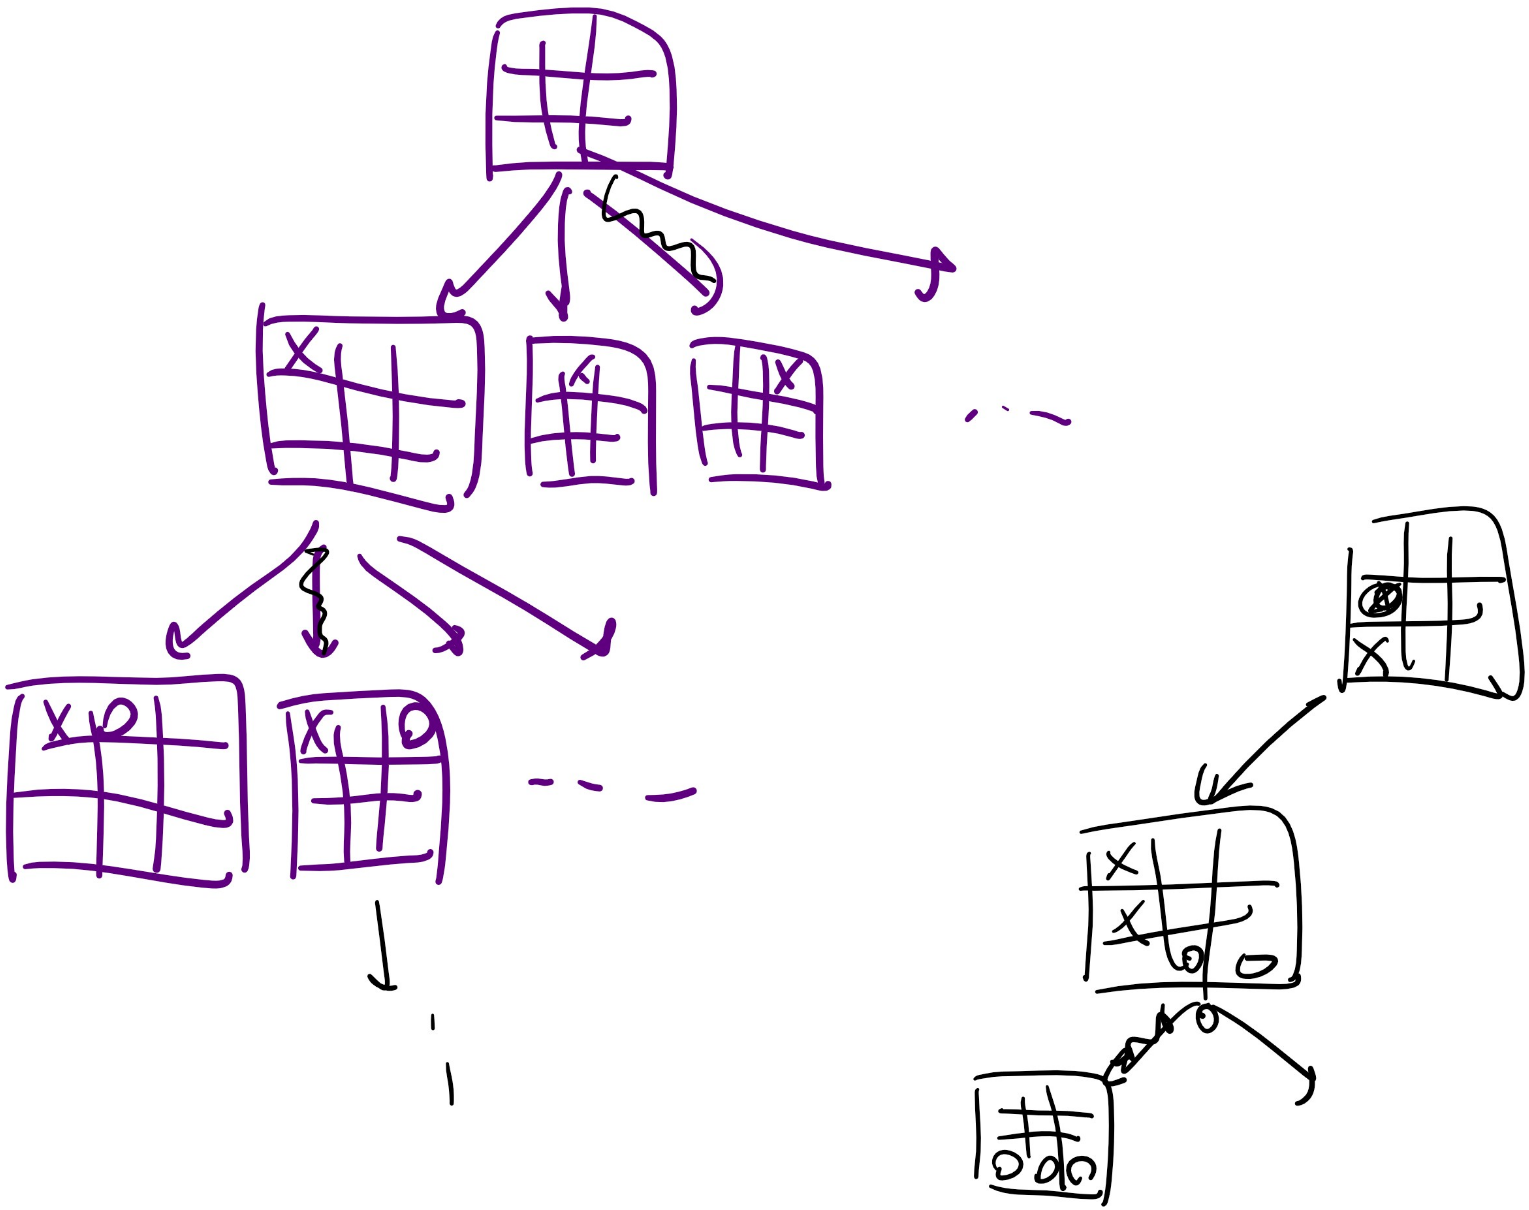

Nim Game Example

Quiz (Graded)

- Ten objects. Pick 1 or 2 each time. Pick the last one to win.
- A: Pick 1.
- B: Pick 2.
- C, D, E: Don't choose.

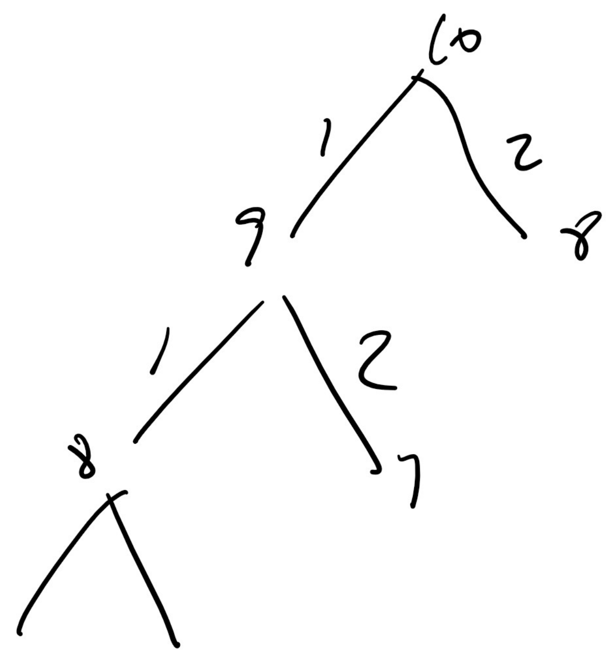

trick
always maintain
multiple of 3.

2 → 1
1 → 2

2 Nim Game Example

Motivation

Minimax Algorithm

Description

- Use DFS on the game tree.

Minimax Algorithm

Description

- Input: a game tree (V, E, c), and the current state s.
- Output: the value of the game at s.
- If s is a terminal state, return $c(s)$.
- If the player is MAX, return the maximum value over all sucessors.

$$\alpha(s) = \max_{s' \in s'(s)} \beta(s')$$

- If the player is MIN, return the minimum value over all sucessors.

$$\beta(s) = \min_{s' \in s'(s)} \alpha(s')$$

Backtracking

Discussion

- The optimal actions (solution paths) can be found by backtracking from all terminal states as in DFS.

$$s^{\star}(s) = \arg\max_{s' \in s'(s)} \beta(s') \text{ for MAX}$$

$$s^{\star}(s) = \arg\min_{s' \in s'(s)} \alpha(s') \text{ for MIN}$$

2 Nim Game Example

Discussion

Minimax Performance

Discussion

- The time and space complexity is the same as DFS. Note that $D = d$ is the maximum depth of the terminal states.

$$T = b + b^2 + ... + b^d$$
$$S = (b - 1) \cdot d$$

Non-deterministic Game

Discussion

- For non-deterministic games in which chance can make a move (dice roll or coin flip), use expected reward or cost instead.

Game Tree with Chance Example

Quiz (Graded)

- Fall 2005 Midterm Q7
- Max can pick L or R. If Max picks L, Chance picks L with probability 0.3 and R with probability 0.7. If Chance picks L, Min picks L to get 3, R to get 2, and if Chance picks R, Min gets 7. If Max picks R, Min picks L to get -1 and R to get 2. What is the value of the game?
- A: -1
- B: 2
- C: 5.5
- D: 5.8
- E: 7

Improvements

Discussion

- Time complexity is a problem because the computer usually has a limited amount of time to "think" and make a move.
- It is possible to reduce the time complexity by removing the branches that will not lead the current player to win. It is called the Alpha-Beta pruning.
- It is also possible to use greedy search with heuristic function (called static board evaluation).