
CS 232:

Artificial Intelligence

Fall 2026

Prof. Carolyn Anderson
Wellesley College

Reminders


- ♦ I have help hours today (W423) and Monday from 3-4 (W323)
- ♦ Tutor hours:
 - Aeka: Thursdays 7:30-8:30 in H103
 - Amber: Saturday 11-12 in Whiteboard Alley
 - Coco: Mondays 7-8 in Whiteboard Alley

THURSDAY, OCTOBER 8
12:45PM TO 2PM
H105



<https://bit.ly/3V2EJLd>

CS Colloquium:
**Tech Ethics and
College Chaplaincy:
Lessons from MIT
Radius**



Recap

Minimax Summary

- ✦ Rank final game states by their final scores (for tic-tac-toe or chess: win, draw, loss).
- ✦ Rank intermediate game states by whose turn it is and the available moves.
 - If it's X's turn, set the rank to that of the *maximum* move available. If a move will result in a win, X should take it.
 - If it's O's turn, set the rank to that of the *minimum* move available. If a move will result in a loss, X should avoid it.

Alpha-Beta Pruning Exercise

Key idea: give up on paths when you realize that they are worse than options you've already explore.

- ♦ Track the **max** score possible for the **minimizing** player (**beta**)
- ♦ Track the **min** score possible for the **maximizing** player (**alpha**)

Whenever the **maximum score for beta** becomes less than the **minimum score for alpha**, the maximizing player can stop searching down this path, because it will never be reached.

AI Needs You Discussion

- ◆ What did you think about this chapter so far?
- ◆ What surprised you the most?

Uncertain Outcomes

Stochastic Transition Model

So far we have considered search problems where the outcomes of our actions was known.

Our transition function took our current state and an action, and told us which state we would end up in next:

$$T(s,a) \rightarrow s'$$

But what if we don't know what will happen next?

Stochastic Transition Function:

$$T(s, a, s') \rightarrow p(s' | s, a)$$

$T(\text{current-state, action, successor/child}) \rightarrow$ probability of reaching child

Expectimax Search


In life, outcomes are often **uncertain**. It can be hard to predict exactly what will happen when we take actions.

- Explicit randomness: rolling dice
- Unpredictable opponents
- Actions can fail: robot might slip while navigating

We can model this using Chance nodes!


$$A = 0.5 \cdot 10 + 0.5 \cdot 10 = 10$$

$$B = 0.9 \cdot 9 + 0.1 \cdot 100 = 8.1 + 10 = 18.1$$



Expectimax Search

Like Minimax, but now there are chance nodes. For chance nodes, we take an **average** of their outcomes (children) **weighted** by the probabilities of each path.



Stochastic Transitions


Markov Decision Processes

- ◆ Known state space
- ◆ Known transition function
- ◆ Transition function is stochastic:

$$T(s, a, s') \rightarrow p(s' | s, a)$$

- ◆ Reward function that tells us the value of being in each state

Markovian transitions depend only on the current state





Escape to the Surface

Suppose we have a fully-observable 4x3 environment with goal states. Ariel begins in the start state and picks an action at each time step.

Actions: *Up, Down, Left, Right*

The game terminates when it reaches a goal state (+1 or -1).





Escape to the Surface

For the action sequence *Up, Up, Right, Right, Right*, what is the probability that Ariel reaches her goal (the surface)?

What are the routes to the goal?

$$\begin{aligned}
 p(\text{surface} | \text{UURRR}, (1,1)) &= \\
 0.8 * 0.8 * 0.8 * 0.8 * 0.8 &= 0.32768 \\
 = 0.1 * 0.1 * 0.1 * 0.1 * 0.8 &= 0.00008
 \end{aligned}$$

Transition Model:



Escape to the Surface

For the action sequence *Up, Up, Right, Right, Right*, what is the probability that Ariel reaches her goal (the surface)?

What are the routes to the goal?

$$p(\text{surface} | \text{UURRR}, (1,1)) =$$

$$p(\text{surface} | \text{all going well})$$

$$+ p(\text{surface} | \text{mistakes work out})$$



$$0.328 + 0.00008$$

Transition Model:

$$p(\text{success}) = 0.8$$

$$p(\text{error to one side}) = 0.1$$


$$(error to other side) = 0.$$





Escape to the Surface

For the action sequence *Up, Up, Right, Right, Right*, what is the probability that Ariel reaches her goal (the surface)?

What are the routes to the goal?



Transition Model:



Solution → Policy

It's harder to plan now that our transition function is stochastic. Instead, we should come up with a **policy**: a function that tells us what to do in every situation.

$$\pi(s) \rightarrow a$$

If the same policy is enacted multiple times, it may lead to different rewards. We care about the **expected utility of a policy**.

Rewards Are Uncertain

When outcomes are uncertain, future rewards are uncertain too. After all, we may never reach them!

For this reason, we typically **discount future rewards**: we set a **weight decay** that decreases the value of rewards based on how far out they are in the future.

