## AI: Representation and Problem Solving Markov Decision Processes



Instructors: Nihar Shah and Tuomas Sandholm
Slide credits: CMU AI and ai.berkeley.edu

The relationship between actions \& consequences is not deterministic, but is stochastic (random)

## Example: Grid World

- A maze-like problem
- The agent lives in a grid
- Walls block the agent's path
- Noisy movement: actions do not always go as planned
- $80 \%$ of the time, the action North takes the agent North (if there is no wall there)
- $10 \%$ of the time, North takes the agent West; $10 \%$ East
- If there is a wall in the direction the agent would have been taken, the agent stays put
- The agent receives rewards each time step

- Reward depends on agent's state and action
- Can be positive, zero or negative
- Goal: maximize sum of rewards


## Grid World Actions

Deterministic Grid World


## Markov Decision Processes

An MDP is defined by:

- A set of states $s \in S$
- A set of actions $a \in A$
- A transition function T(s, a, s')
- Probability that a from s leads to $\mathrm{s}^{\prime}$, i.e., $\mathrm{P}\left(\mathrm{s}^{\prime} \mid \mathrm{s}, \mathrm{a}\right)$
- A reward function $R\left(s, a, s^{\prime}\right)$

All of this information is known beforehand


## What is Markov about MDPs?

"Markov" generally means that given the present state, the future and the past are independent

For Markov decision processes, "Markov" means action outcomes depend only on the current state

$$
\begin{aligned}
& P\left(S_{t+1}=s^{\prime} \mid S_{t}=s_{t}, A_{t}=a_{t}, S_{t-1}=s_{t-1}, A_{t-1}, \ldots S_{0}=s_{0}\right) \\
& \quad= \\
& P\left(S_{t+1}=s^{\prime} \mid S_{t}=s_{t}, A_{t}=a_{t}\right)
\end{aligned}
$$



Andrey Markov (1856-1922)

## Goal: Maximize sum of rewards

## Finite time horizon setting:

- Consider MDP for T time steps
- Rewards summed over the time steps
- Want to choose actions that maximize expected reward
$\mathrm{E}\left[\sum_{t=1}^{T}\right.$ Reward at time t]


## Markov Decision Processes

What actions will yield the highest expected reward?


Suppose you choose the action "up". Then suppose:

- There is an $80 \%$ chance you go up, and get a reward of +2
- There is a $10 \%$ chance you go left, and get a reward of -1
- There is a $10 \%$ chance you go right, and get a reward of -1
- The expected reward in this step under this action is $0.8^{*}(+2)+0.1^{*}(-1)+0.1^{*}(-1)=1.4$

Want to choose actions that maximize total reward across many time steps (not just in the current time step)

## MDP Search Trees



## Policies

In deterministic single-agent search problems, we wanted an optimal plan, or sequence of actions, from start to a goal

For MDPs, we want an optimal policy

- A policy $\pi$ gives an action for each state
- An optimal policy is one that maximizes expected utility if followed



## Solution method 1: Expectimax search

Suppose the MDP has only 1 state and will run only for 1 time step. If you were to take Left, what is the expected reward?
A) 12
B) 8
C) 7
D) 4


Suppose the MDP has only 1 state and will run only for 1 time step. If you were to take Center, what is the expected reward?
A) 12
B) 8
C) 7
D) 4


Suppose the MDP has only 1 state and will run only for 1 time step. If you were to take Right, what is the expected reward?
A) 12
B) 8
C) 7
D) 4


Suppose the MDP has only 1 state and will run only for 1 time step. Which action should we choose?
A) Left
B) Center
C) Right
D) Cannot determine


## Solution method 1: Expectimax search

- Starting from the start state, expand the tree to depth $k$, for some chosen $k$
- Previous example was for depth 1
- Here is an example for depth 2...



## Solution method 1: Expectimax search

- Starting from the start state, expand the tree to depth $k$, for some chosen $k$
- Move up the tree, recursively computing the reward obtained
- Compute expectations at $\bigcirc$ nodes and max at $\Delta$ nodes

function EXPECTIMAX(node, depth):
if depth $=0$ :
return 0
if node is a MAX node (denoted as s):
value $=-\infty$
for each possible action a:
value $=\max ($ value $\operatorname{EXPECTIMAX}(\mathrm{s}, \mathrm{a})$, depth -1$))$
return value
else if node is an EXPECTATION node $\bigcirc$ (denoted as $(s, a)$ ):
value $=0$
for each child of node $s^{\prime}$, with probability given by the transition $T\left(s, a, s^{\prime}\right)$ :
value $+=T\left(s, a, s^{\prime}\right)^{*}\left(R\left(s, a, s^{\prime}\right)+\right.$ EXPECTIMAX(s', depth $\left.)\right)$
return value


## Goal: Maximize sum of rewards

## Finite time horizon setting:

- Consider MDP for T time steps
- Rewards summed over the time steps
- Want to choose actions that maximize expected reward

$$
\mathrm{E}\left[\sum_{t=1}^{T} \text { Reward at time } \mathrm{t}\right]
$$

Will the optimal choice of actions depend on time?
If you are in a state $s$ at time 1 versus in the state $s$ at time 10,

- can the optimal action be different?




## Goal: Maximize sum of rewards

## Finite time horizon setting:

- Consider MDP for T time steps
- Rewards summed over the time steps
- Want to choose actions that maximize expected reward

$$
\mathrm{E}\left[\sum_{t=1}^{T} \text { Reward at time } \mathrm{t}\right]
$$

Will the optimal choice of actions depend on time?
If you are in a state $s$ at time 1 versus in the state $s$ at time 10, YES

- can the optimal action be different?


## Goal: Maximize sum of rewards

## Finite time horizon setting:

- Consider MDP for T time steps
- Rewards summed over the time steps
- Want to choose actions that maximize expected reward

$$
\mathrm{E}\left[\sum_{t=1}^{T} \text { Reward at time } \mathrm{t}\right]
$$

## Challenges

- Optimal action depends on time step
- Need to compute it separately for each time step...
- Computational cost + headache
- May not know horizon T in advance


## Goal: Maximize sum of rewards

 Infinite time horizon setting:- Want to choose actions that maximize expected reward $\mathrm{E}\left[\sum_{t=1}^{\infty}\right.$ Reward at time t]


## Goal: Maximize sum of rewards

## Infinite time horizon setting:

- Want to choose actions that maximize expected reward
$\mathrm{E}\left[\sum_{t=1}^{\infty}\right.$ Reward at time t]
Challenge
This can be infinite for many choices of actions...
Infinity vs. infinity?


## Discounting

It's reasonable to maximize the sum of rewards
It's also reasonable to prefer rewards now to rewards later
One solution: values of rewards decay exponentially


## Goal: Maximize sum of rewards

Infinite time horizon setting with discounted rewards:
Want to choose actions that maximize expected reward
$\mathrm{E}\left[\sum_{t=1}^{\infty} \gamma^{t}\right.$ Reward at time t]
for some $\gamma$ in $(0,1)$

If you are in some state $s$ at some time $t_{1}$ versus
if you are in that state $s$ at some time $t_{2}$

## Goal: Maximize sum of rewards

$$
\mathrm{E}\left[\sum_{t=1}^{\infty} \gamma^{t} \text { Reward at time } \mathrm{t}\right]
$$

Does the optimal action need to differ with time?

1. Markov property means that the MDP evolution does not change with time
2. What about optimal actions? Let's look at future reward starting at $\mathrm{t}_{1}$ or $\mathrm{t}_{2}$

Finite horizon: Future reward $=\mathrm{E}\left[\sum_{t=\mathrm{t}_{1}}^{T}\right.$ Reward at time t$]$ vs. $\mathrm{E}\left[\sum_{t=\mathrm{t}_{2}}^{T}\right.$ Reward at time t$]$ Infinite horizon with discounted reward:
$\mathrm{E}\left[\sum_{t=\mathrm{t}_{1}}^{\infty} \gamma^{t}\right.$ Reward at time t$]=\mathrm{E}\left[\sum_{t=1}^{\infty} \gamma^{t}\right.$ Reward at time t$]$ (via change of variables) vs. $\mathrm{E}\left[\sum_{t=\mathrm{t}_{2}}^{\infty} \gamma^{t}\right.$ Reward at time t$]=\mathrm{E}\left[\sum_{t=1}^{\infty} \gamma^{t}\right.$ Reward at time t$]$

Time invariant!!

## Policies

In deterministic single-agent search problems, we wanted an optimal plan, or sequence of actions, from start to a goal

For MDPs, we want an optimal policy

- A policy $\pi$ gives an action for each state
- An optimal policy is one that maximizes expected utility if followed

- Action depends only on the state you are in, and not on when
- Policy $\pi: S \rightarrow A$


## Solution method 2: Value iteration

## Optimal Quantities

- The value (utility) of a state s:
$V^{*}(s)=$ expected total reward starting in $s$ and acting optimally
- The value (utility) of a $q$-state ( $s, a$ ):
$Q^{*}(s, a)=$ expected total reward starting out having taken action a from state $s$ and (thereafter) acting optimally
- The optimal policy:
$\pi^{*}(s)=$ optimal action from state $s$


## Optimal Quantities

## Bellman equations

- The value (utility) of a state s:
$\mathrm{V}^{*}(\mathrm{~s})=$ expected total reward starting in $s$ and acting optimally

$$
V^{*}(s)=\max _{a} \sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V^{*}\left(s^{\prime}\right)\right]
$$ for every state s

- The value (utility) of a qstate ( $\mathrm{s}, \mathrm{a}$ ):
$Q^{*}(s, a)=$ expected total reward starting out having taken action a from state $s$ and (thereafter) acting optimally
- The optimal policy: $\pi^{*}(\mathrm{~s})=$ optimal action from state s


## Optimal Quantities

- The value (utility) of a state s:
$\mathrm{V}^{*}(\mathrm{~s})=$ expected total reward starting in $s$ and acting optimally
- The value (utility) of a qstate ( $\mathrm{s}, \mathrm{a}$ ):
$Q^{*}(s, a)=$ expected total reward starting out having taken action a from state $s$ and (thereafter) acting optimally
- The optimal policy: $\pi^{*}(\mathrm{~s})=$ optimal action from state s


## Question: If you have $Q^{*}$, how to get $\mathrm{V}^{*}$ ?

$$
V^{*}(s)=\max _{a} Q^{*}(s, a)
$$

## Optimal Quantities

- The value (utility) of a state s:
$\mathrm{V}^{*}(\mathrm{~s})=$ expected total reward starting in $s$ and acting optimally
- The value (utility) of a qstate ( $\mathrm{s}, \mathrm{a}$ ):
$Q^{*}(s, a)=$ expected total reward starting out having taken action a from state s and (thereafter) acting optimally
- The optimal policy: $\pi^{*}(\mathrm{~s})=$ optimal action from state s


## Question: If you have $\mathrm{Q}^{*}$, how to get $\pi^{*}$ ?

$$
\pi^{*}(\mathrm{~s})=\underset{a}{\arg \max _{a}} Q^{*}(s, a)
$$

## Optimal Quantities

- The value (utility) of a state s:
$\mathrm{V}^{*}(\mathrm{~s})=$ expected total reward starting in $s$ and acting optimally
- The value (utility) of a qstate ( $\mathrm{s}, \mathrm{a}$ ):
$Q^{*}(s, a)=$ expected total reward starting out having taken action a from state $s$ and (thereafter) acting optimally
- The optimal policy: $\pi^{*}(\mathrm{~s})=$ optimal action from state s


## Question: If you have $\mathrm{V}^{*}$, how to get $\mathrm{Q}^{*}$ ?

$$
Q^{*}(s, a)=\sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V^{*}\left(s^{\prime}\right)\right]
$$

## Optimal Quantities

- The value (utility) of a state s:
$\mathrm{V}^{*}(\mathrm{~s})=$ expected total reward starting in $s$ and acting optimally
- The value (utility) of a qstate ( $\mathrm{s}, \mathrm{a}$ ):
$Q^{*}(s, a)=$ expected total reward starting out having taken action a from state $s$ and (thereafter) acting optimally
- The optimal policy: $\pi^{*}(\mathrm{~s})=$ optimal action from state s


## Question: If you have $\mathrm{V}^{*}$, how to get $\pi^{*}$ ?

$$
\pi^{*}(s)=\arg \max _{a} \sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V^{*}\left(s^{\prime}\right)\right]
$$

"Policy extraction"

## Solution method 2: Value iteration

$V^{*}(s)=\max _{a} \sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V^{*}\left(s^{\prime}\right)\right]$
If we can estimate $\mathrm{V}^{*}$, we can find the optimal policy

Key idea:

- Initialize $\mathrm{V}^{*}$ (either a guess or set it to 0 )
- Keep updating the guess using the equation above


## Solution method 2: Value iteration

Start with initial guess of the value function $\mathrm{V}, \mathrm{e} . \mathrm{g} ., \mathrm{V}(\mathrm{s})=0$ for all s
Given vector of $\mathrm{V}(\mathrm{s})$ values, do one iteration of expectimax for each state:

$$
V_{k+1}(s) \leftarrow \max _{a} \sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V_{k}\left(s^{\prime}\right)\right]
$$

Repeat

## Solution method 2: Value iteration

Example on the board

## Solution method 2: Value iteration

Start with initial guess of the value function $V$, e.g., $V(s)=0$ for all $s$
Given vector of $\mathrm{V}(\mathrm{s})$ values, do one iteration of expectimax for each state:

$$
V_{k+1}(s) \leftarrow \max _{a} \sum_{s^{\prime}} T\left(s, a, s^{\prime}\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V_{k}\left(s^{\prime}\right)\right]
$$

Repeat

Complexity of each iteration:

$$
O\left(S^{2} A\right)
$$

Theorem: will converge to unique optimal values

## Convergence

How do we know the $\mathrm{V}_{\mathrm{k}}$ vectors are going to converge?

Proof sketch: Suppose rewards are bounded in [ $R_{\text {MIN }}, R_{\text {MAX }}$ ]

- For any state $\mathrm{V}_{\mathrm{k}}$ and $\mathrm{V}_{\mathrm{k}+1}$ can be viewed as depth $\mathrm{k}+1$ expectimax results in nearly identical search trees
- The difference is that on the bottom layer, $\mathrm{V}_{\mathrm{k}+1}$ has actual rewards while $\mathrm{V}_{\mathrm{k}}$ has zeros
- That last layer is at best all $R_{\text {MAX }}$
- It is at worst $\mathrm{R}_{\text {MIN }}$
- But everything is discounted by $\gamma^{k}$ that far out
- So $\mathrm{V}_{\mathrm{k}}$ and $\mathrm{V}_{\mathrm{k}+1}$ are at most $\gamma^{k} \max |\mathrm{R}|$ different
- So as $k$ increases, the values converge


