## Warm-up as You Walk In

## Given

- Set actions (persistent/static)
- Set states (persistent/static)
- Function T(s,a,s_prime)

Write the pseudo code for:

- function $V(\underline{s})$ return value
that implements:

$$
V(s)=\max _{a \in \text { actions }} \sum_{s^{\prime} \in \text { states }} T\left(s, a, s^{\prime}\right) V\left(s^{\prime}\right)
$$

## Announcements

## Assignments:

- P3
- Due Thu 3/7, 10 pm


## Spring Break!

- No recitation this Friday
- HW7 (online): out Wed 3/6, due Tue 3/19
- P4: out after break, due Thu $3 / 28$


## AI: Representation and Problem Solving Markov Decision Processes



Instructors: Pat Virtue \& Stephanie Rosenthal
Slide credits: CMU AI and http://ai.berkeley.edu

## Minimax Notation



$$
V(s)=\underset{\substack{a \\ \text { where } \\ s^{\prime}}}{\max ^{\prime} V\left(s^{\prime}\right), \quad A a_{2}} \rightarrow C
$$

$$
\begin{aligned}
& \hat{a}=\underset{a}{\operatorname{argmax}} V\left(s^{\prime}\right), \\
& \text { where } s^{\prime}=\operatorname{result}(s, a)
\end{aligned}
$$

## Expectations



Max node notation

$$
\begin{aligned}
& V(s)=\max _{a} V\left(s^{\prime}\right) \\
& \quad \text { where } s^{\prime}=\operatorname{result}(s, a)
\end{aligned}
$$

Chance node notation

$$
V(s)=\underline{\underline{\sum_{\underline{s^{\prime}}}}} P\left(s^{\prime}\right) V\left(s^{\prime}\right)
$$

## Piazza Poll 1

Expectimax tree search:
Which action do we choose?
A) Left
B) Center
C) Right
D) Eight


## Piazza Poll 1

$$
\hat{a}=\underset{a}{\operatorname{argmax}} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right) V\left(s^{\prime}\right)
$$

Expectimax tree search:
Which action do we choose?
A) Left
B) Center
C) Right
D) Eight


## Expectimax Notation

$$
V(s)=\max _{a} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right) V\left(s^{\prime}\right)
$$

## Warm-up as You Walk In

## Given

- Set actions (persistent/static)
- Set states (persistent/static)

$$
P\left(s^{\prime} \mid s, a\right)
$$

- Function $T$ (s,a,s_prime)

Write the pseudo code for:

- function $V(s)$ return value
that implements:

$$
V(s)=\max _{a \in \text { actions }} \sum_{s^{\prime} \in \text { states }} T\left(s, a, s^{\prime}\right) V\left(s^{\prime}\right)
$$



## MDP Notation

| Standard expectimax: | $V(s)=\max _{a} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right) V\left(s^{\prime}\right)$ |
| :--- | :--- |
| Bellman equations: | $V(s)=\max _{a} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V\left(s^{\prime}\right)\right]$ |
| Value iteration: | $V_{k+1}(s)=\max _{a} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V_{k}\left(s^{\prime}\right)\right], \quad \forall s$ |
| Q-iteration: | $Q_{k+1}(s, a)=\sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \max _{a^{\prime}} Q_{k}\left(s^{\prime}, a^{\prime}\right)\right], \quad \forall s, a$ |
| Policy extraction: | $\pi_{V}(s)=\underset{a}{\operatorname{argmax}} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V\left(s^{\prime}\right)\right], \quad \forall s$ |
| Policy evaluation: | $V_{k+1}^{\pi}(s)=\sum_{s^{\prime}} P\left(s^{\prime} \mid s, \pi(s)\right)\left[R\left(s, \pi(s), s^{\prime}\right)+\gamma V_{k}^{\pi}\left(s^{\prime}\right)\right], \quad \forall s$ |
| Policy improvement: | $\pi_{n e w}(s)=\underset{a}{\operatorname{argmax}} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V^{\left.\pi_{o l d}\left(s^{\prime}\right)\right], \quad \forall s}\right.$ |

## MDP Notation

$\begin{array}{ll}\text { Standard expectimax: } & V(s)=\max _{a} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right) V\left(s^{\prime}\right) \\ \text { Bellman equations: } & V(s)=\max _{a} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V\left(s^{\prime}\right)\right] \\ \text { Value iteration: } & V_{k+1}(s)=\max _{a} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V_{k}\left(s^{\prime}\right)\right], \quad \forall s \\ \text { Q-iteration: } & Q_{k+1}(s, a)=\sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \max _{a^{\prime}} Q_{k}\left(s^{\prime}, a^{\prime}\right)\right], \quad \forall s, a \\ \text { Policy extraction: } & \pi_{V}(s)=\underset{a}{\operatorname{argmax}} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V\left(s^{\prime}\right)\right], \quad \forall s \\ \text { Policy evaluation: } & V_{k+1}^{\pi}(s)=\sum_{s^{\prime}} P\left(s^{\prime} \mid s, \pi(s)\right)\left[R\left(s, \pi(s), s^{\prime}\right)+\gamma V_{k}^{\pi}\left(s^{\prime}\right)\right], \quad \forall s \\ \text { Policy improvement: } & \pi_{n e w}(s)=\underset{a}{\operatorname{argmax}} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V^{\left.\pi_{o l d}\left(s^{\prime}\right)\right], \quad \forall s}\right.\end{array}$

Non-Deterministic Search


## 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
- Small "living" reward each step (can be negative)
- Big rewards come at the end (good or bad)
- Goal: maximize sum of rewards


## Grid World Actions

Deterministic Grid World


Stochastic 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 s', i.e., $\mathrm{P}\left(\mathrm{s}^{\prime} \mid \mathrm{s}, \mathrm{a}\right)$
- Also called the model or the dynamics
- A reward function $\mathrm{R}\left(\mathrm{s}, \mathrm{a}, \mathrm{s}^{\prime}\right)$
- Sometimes just R(s) or R(s')
- A start state
- Maybe a terminal state


MDPs are non-deterministic search problems

- One way to solve them is with expectimax search
- We'll have a new tool soon


## Demo of Gridworld

## 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

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



Andrey Markov (1856-1922)

This is just like search, where the successor function could only depend on the current state (not the history)

## 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 $\pi^{*}: \mathrm{S} \rightarrow \mathrm{A}$

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


## Expectimax didn't compute entire policies



Optimal policy when $R\left(s, a, s^{\prime}\right)=-0.03$ for all non-terminals s

- It computed the action for a single state only


## Piazza Poll 2

Which sequence of optimal policies matches the following sequence of living rewards: $\{-0.01,-0.03,-0.04,-2.0\}$
I. $\{A, B, C, D\}$
II. $\{B, C, A, D\}$
III. $\{D, C, B, A\}$
IV. $\{D, A, C, B\}$

B)

C)

D)

| - | - | - | $\square$ |
| :--- | :--- | :--- | :--- |
| 1 |  | - | $\square$ |
| - | - | - | 1 |

## Piazza Poll 2

Which sequence of optimal policies matches the following sequence of living rewards:
$\{-0.01,-0.03,-0.04,-2.0\}$


Optimal Policies

$R(s)=-0.01$

$R(s)=-0.4$

$R(s)=-0.03$

$R(s)=-2.0$

## Example: Racing



## Example: Racing

A robot car wants to travel far, quickly
Three states: Cool, Warm, Overheated
Two actions: Slow, Fast
Going faster gets double reward


Overheated

## Racing Search Tree



## MDP Search Trees

Each MDP state projects an expectimax-like search tree


Utilities of Sequences


## Utilities of Sequences

What preferences should an agent have over reward sequences?

More or less? $\quad[1,2,2] \quad$ or $\quad[2,3,4]$

Now or later? $\quad[0,0,1]$ or $[1,0,0]$


## 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


## Discounting

## How to discount?

- Each time we descend a level, we multiply in the discount once


## Why discount?

- Sooner rewards probably do have higher utility than later rewards
- Also helps our algorithms converge



## Piazza Poll 3

What is the value of $\mathrm{U}[2,4,8]$ with $\gamma=0.5$ ?
A. 3
B. 6
C. 7
D. 14

Bonus: What is the value of $\mathrm{U}[8,4,2]$ with $\gamma=0.5$ ?

Piazza Poll 3
What is the value of $U[2,4,8]$ with $\gamma=0.5$ ?
A. 3
B. 6
C. 7
D. 14

$$
\gamma^{0} 2+\gamma^{\prime 4}+\gamma^{2} 8
$$

$$
1 \cdot 2+0.54+0.25 \cdot 8
$$

$2+2$
Bonus: What is the value of $\mathrm{U}[8,4,2]$ with $\gamma=0.5$ ?

$$
10,5
$$

## Stationary Preferences

Theorem: if we assume stationary preferences:

$$
\int \quad\left[a_{1}, a_{2}, \ldots\right] \succ\left[b_{1}, b_{2}, \ldots\right]
$$

$$
\Uparrow
$$

$$
\left[r, a_{1}, a_{2}, \ldots\right] \succ\left[r, b_{1}, b_{2}, \ldots\right]
$$



Then: there are only two ways to define utilities

- Additive utility:

$$
U\left(\left[r_{0}, r_{1}, r_{2}, \ldots\right]\right)=r_{0}+r_{1}+r_{2}+\cdots
$$

- Discounted utility:

$$
U\left(\left[r_{0}, r_{1}, r_{2}, \ldots\right]\right)=r_{0}+\underline{\gamma} \underline{r}_{1}+\underline{\gamma}^{2} r_{2} \cdots
$$

## Discounting

Given:

| 10 |  |  |  | 1 |
| :---: | :---: | :---: | :---: | :---: |
| a | b | c | d | e |

- Actions: East, West, and Exit (only available in exit states a, e)
- Transitions: deterministic

For $\gamma=1$, what is the optimal policy?


For $\gamma=0.1$, what is the optimal policy?
For which $\gamma$ are West and East equally good when in state d?


## Infinite Utilities?!

- Problem: What if the game lasts forever? Do we get infinite rewards?


## - Solutions:

- Finite horizon: (similar to depth-limited search)
- Terminate episodes after a fixed T steps (e.g. life)
- Gives nonstationary policies ( $\pi$ depends on time left)
- Discounting: use $0<\gamma<1$


$$
U(\underbrace{\left.R_{m}, \ldots r_{\infty}\right]})=\sum_{t=0}^{\infty} \gamma^{t} r_{t} \leq R_{\max } /(1-\gamma)
$$

- Smaller $\gamma$ means smaller "horizon" - shorter term focus

- Absorbing state: guarantee that for every policy, a terminal state will eventually be reached (like "overheated" for racing)


## Recap: Defining MDPs

Markov decision processes:

- Set of states S
- Start state $\mathrm{s}_{0}$
- Set of actions A
- Transitions P(s'|s,a)(or T(s,a, s'))
- Rewards $\frac{R\left(s, a, s^{\prime}\right)}{\uparrow \uparrow \uparrow}$ (and discount $\gamma$ )


MDP quantities so far:

- Policy=Choice of @ction)for each state
- Utility = sum of (discounted) rewards


Solving MDPs


## Optimal Quantities

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

- The optimal policy:
$\pi^{*}(s)=$ optimal action from state $s$


## Snapshot of Demo - Gridworld V Values



VALUES AFTER 100 ITERATIONS

Noise $=0.2$
Discount $=0.9$
Living reward $=0$

## Snapshot of Demo - Gridworld Q Values

## Gridworld Display



Q-VALUES AFTER 100 ITERATIONS

Noise $=0.2$
Discount = 0.9
Living reward = 0

## Values of States

Fundamental operation: compute the (expectimax) value of a state

- Expected utility under optimal action
- Average sum of (discounted) rewards
- This is just what expectimax computed!

Recursive definition of value:

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

## Racing Search Tree



## Racing Search Tree



## Racing Search Tree

We're doing way too much work with expectimax!

Problem: States are repeated

- Idea: Only compute needed quantities once

Problem: Tree goes on forever

- Idea: Do a depth-limited computation, but with increasing depths until change is small
- Note: deep parts of the tree eventually don't matter if $\gamma<1$



## Time-Limited Values

Key idea: time-limited values
Define $V_{k}(s)$ to be the optimal value of $s$ if the game ends in $k$ more time steps

- Equivalently, it's what a depth-k expectimax would give from s



## $\mathrm{k}=0$

| $\Delta$ | $\Delta$ | $\Delta$ | $\square$ |
| :---: | :---: | :---: | :---: |
| 0.00 | 0.00 | 0.00 | 0.00 |
| 0.00 |  | 0.00 | 0.00 |
| 0.00 | 0.00 | 0.00 | 0.00 |
|  |  |  |  |

VALUES AFTER 0 ITERATIONS

Noise $=0.2$
Discount $=0.9$
Living reward = 0
$\mathrm{k}=1$

| $\Delta$ | $\Delta$ |  |  |
| :---: | :---: | :---: | :---: |
| 0.00 | 0.00 | 0.00 | 1.00 |
| $\Delta$ |  |  |  |
| 0.00 |  | 0.00 | -1.00 |
| 0.00 | 0.00 | 0.00 | 0.00 |

VALUES AFTER 1 ITERATIONS

Noise $=0.2$
Discount $=0.9$
Living reward $=0$
$\mathrm{k}=2$

|  |  |  |  |
| :---: | :---: | :---: | :---: |
| 0.00 | 0.00 | 0.72 | 1.00 |
| 0.00 |  | 0.00 | -1.00 |
| 0.00 | 0.00 | 0.00 | 0.00 |
| 0 |  |  | - |

VALUES AFTER 2 ITERATIONS

Noise $=0.2$
Discount $=0.9$
Living reward $=0$

$$
k=3
$$

| 0.00 | 0.52, | 0.78 | 1.00 |
| :---: | :---: | :---: | :---: |
| 0 |  | - | $\square$ |
| 0.00 |  | 0.43 | -1.00 |
| 0.00 | 0.00 | 0.00 | 0.00 |
|  |  |  | - |

VALUES AFTER 3 ITERATIONS

Noise $=0.2$
Discount $=0.9$
$\mathrm{k}=4$


VALUES AFTER 4 ITERARIONS

Noise $=0.2$
Discount $=0.9$
Living reward $=0$
$\mathrm{k}=5$


VALUES AFTER 5 ITERATIONS

Noise $=0.2$
Discount $=0.9$
Living reward $=0$
$\mathrm{k}=6$


VALUES AFTER 6 ITERATIONS

Noise $=0.2$
Discount $=0.9$
Living reward $=0$
$k=7$


VALUES AFTER 7 ITERATIONS

Noise $=0.2$
Discount $=0.9$
Living reward $=0$
$\mathrm{k}=8$


VALUES AFTER 8 ITERATIONS

Noise $=0.2$
Discount $=0.9$
Living reward $=0$
$\mathrm{k}=9$


VALUES AFTER 9 ITERARIONS

Noise $=0.2$
Discount $=0.9$
Living reward $=0$
$\mathrm{k}=10$


VALUES AFTEER 10 ITERATIONS

Noise $=0.2$
Discount $=0.9$
Living reward $=0$
$\mathrm{k}=11$


VALUES AFTEER 11 ITERATIONS

Noise $=0.2$
Discount $=0.9$
Living reward $=0$
$\mathrm{k}=12$

| 0.64 | 0.74 | 0.85 | 1.00 |
| :---: | :---: | :---: | :---: |
| - |  | - |  |
| 0.57 |  | 0.57 | -1.00 |
| - |  | $\triangle$ |  |
| 0.49 | 40.42 | 0.47 | 40.28 |

VALUES AFIER 12 IMERATIONS

Noise $=0.2$
Discount $=0.9$
Living reward $=0$
$\mathrm{k}=100$


VALUES AFTER 100 ITERATIONS

Noise $=0.2$
Discount $=0.9$
Living reward $=0$

## Computing Time-Limited Values



## Value Iteration



## Value Iteration

Start with $\mathrm{V}_{0}(\mathrm{~s})=0$ : no time steps left means an expected reward sum of zero
Given vector of $\mathrm{V}_{\mathrm{k}}(\mathrm{s})$ values, do one ply of expectimax from 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 until convergence

Complexity of each iteration: $\mathrm{O}\left(\mathrm{S}^{2} \mathrm{~A}\right)$
Theorem: will converge to unique optimal values

- Basic idea: approximations get refined towards optimal values

- Policy may converge long before values do


## Example: Value Iteration



## Convergence

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

Case 1: If the tree has maximum depth $M$, then $V_{M}$ holds the actual untruncated values

## Case 2: If the discount is less than 1

- Sketch: 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 $\gamma^{\mathrm{k}} \max |\mathrm{R}|$ different
- So as $k$ increases, the values converge


## MDP Notation

| Standard expectimax: | $V(s)=\max _{a} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right) V\left(s^{\prime}\right)$ |
| :--- | :--- |
| Bellman equations: | $V(s)=\max _{a} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V\left(s^{\prime}\right)\right]$ |
| Value iteration: | $V_{k+1}(s)=\max _{a} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V_{k}\left(s^{\prime}\right)\right], \quad \forall s$ |
| Q-iteration: | $Q_{k+1}(s, a)=\sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \max _{a^{\prime}} Q_{k}\left(s^{\prime}, a^{\prime}\right)\right], \quad \forall s, a$ |
| Policy extraction: | $\pi_{V}(s)=\underset{a}{\operatorname{argmax}} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V\left(s^{\prime}\right)\right], \quad \forall s$ |
| Policy evaluation: | $V_{k+1}^{\pi}(s)=\sum_{s^{\prime}} P\left(s^{\prime} \mid s, \pi(s)\right)\left[R\left(s, \pi(s), s^{\prime}\right)+\gamma V_{k}^{\pi}\left(s^{\prime}\right)\right], \quad \forall s$ |
| Policy improvement: | $\pi_{n e w}(s)=\underset{a}{\operatorname{argmax}} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V^{\left.\pi_{o l d}\left(s^{\prime}\right)\right], \quad \forall s}\right.$ |

## MDP Notation

$\begin{array}{ll}\text { Standard expectimax: } & V(s)=\max _{a} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right) V\left(s^{\prime}\right) \\ \text { Bellman equations: } & V(s)=\max _{a} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V\left(s^{\prime}\right)\right] \\ \text { Value iteration: } & V_{k+1}(s)=\max _{a} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V_{k}\left(s^{\prime}\right)\right], \quad \forall s \\ \text { Q-iteration: } & Q_{k+1}(s, a)=\sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma \max _{a^{\prime}} Q_{k}\left(s^{\prime}, a^{\prime}\right)\right], \quad \forall s, a \\ \text { Policy extraction: } & \pi_{V}(s)=\underset{a}{\operatorname{argmax}} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V\left(s^{\prime}\right)\right], \quad \forall s \\ \text { Policy evaluation: } & V_{k+1}^{\pi}(s)=\sum_{s^{\prime}} P\left(s^{\prime} \mid s, \pi(s)\right)\left[R\left(s, \pi(s), s^{\prime}\right)+\gamma V_{k}^{\pi}\left(s^{\prime}\right)\right], \quad \forall s \\ \text { Policy improvement: } & \pi_{n e w}(s)=\underset{a}{\operatorname{argmax}} \sum_{s^{\prime}} P\left(s^{\prime} \mid s, a\right)\left[R\left(s, a, s^{\prime}\right)+\gamma V^{\left.\pi_{o l d}\left(s^{\prime}\right)\right], \quad \forall s}\right.\end{array}$

Next Time: Policy-Based Methods

