## 10-701 <br> Machine Learning

Reinforcement learning (RL)

## Markov decision process (MDP) with actions



## Value computation

- An obvious question for such models is what is combined expected value for each state
- What can we expect to earn over our life time if we become Asst. prof.?
- What if we go to industry?

Before we answer this question, we need to define a model for future rewards:

- The value of a current award is higher than the value of future awards
- Inflation, confidence
- Example: Lottery


## Discounted rewards

- The discounted rewards model is specified using a parameter $\gamma$
- Total rewards = current reward +
$\gamma($ reward at time $t+1)+$
$\gamma^{2}$ (reward at time t+2) +
$\gamma^{\mathrm{k}}$ (reward at time $\left.\mathrm{t}+\mathrm{k}\right)+$
infinite sum


## Discounted awards

- The discounted award model is specified using a parameter $\gamma$
- Total awards = current award + $\gamma$ (award at time $t+1)+$ $v^{2}$ (award at time ++ ? $)+$


## Converges if $0<\gamma<1 \quad$ <) +

infinite sum

## Determining the total rewards in a state

- Define $\mathrm{J}^{*}\left(\mathrm{~s}_{\mathrm{i}}\right)=$ expected discounted sum of rewards when starting at state $\mathrm{s}_{\mathrm{i}}$
- How do we compute $\mathrm{J}^{\star}\left(\mathrm{s}_{\mathrm{i}}\right)$ ?

Factors expected pay
$\begin{aligned} & \\ & *\left(s_{i}\right)=r_{i}+\gamma X \\ & \begin{array}{l}\text { for all possible } \\ \text { transitions for step } i\end{array} \\ &=r_{i}+\gamma\left(p_{i 1} J *\left(s_{1}\right)+p_{i 2} J *\left(s_{2}\right)+\cdots p_{i n} J *\left(s_{n}\right)\right)\end{aligned}$

How can we solve this?

## Iterative approaches

- Solving in closed form is possible, but may be time consuming.
- It also doesn't generalize to non-linear models
- Alternatively, this problem can be solved in an iterative manner
- Lets define $\mathrm{J}^{\mathrm{t}}\left(\mathrm{s}_{\mathrm{i}}\right)$ as the expected discounted rewards after k steps
- How can we compute $\mathrm{J}^{\mathrm{t}}\left(\mathrm{s}_{\mathrm{i}}\right)$ ?

$$
\begin{aligned}
& J^{1}\left(S_{i}\right)=r_{i} \\
& J^{2}\left(S_{i}\right)=r_{i}+\gamma\left(\sum_{k} p_{i, k} J^{1}\left(s_{k}\right)\right) \\
& J^{t+1}\left(S_{i}\right)=r_{i}+\gamma\left(\sum_{k} p_{i, k} J^{t}\left(s_{k}\right)\right)
\end{aligned}
$$

## Iterative approaches

- We know how to solve this!
. Lets fill the dynamic programming table
- Lets detine $\mathrm{J}^{\mathrm{K}}\left(\mathrm{s}_{\mathrm{i}}\right)$ as the expected discounted awards atter k steps
- But wait ...

This is a never ending task!

$$
\begin{aligned}
& J^{2}\left(S_{i}\right)=r_{i}+\gamma\left(\sum_{k} p_{i, k} J^{1}\left(s_{k}\right)\right) \\
& J^{t+1}\left(S_{i}\right)=r_{i}+\gamma\left(\sum_{k} p_{i, k} J^{t}\left(s_{k}\right)\right)
\end{aligned}
$$

## When do we stop?

$$
\begin{aligned}
& J^{1}\left(S_{i}\right)=r_{i} \\
& J^{2}\left(S_{i}\right)=r_{i}+\gamma\left(\sum_{k} p_{i, k} J^{1}\left(s_{k}\right)\right) \\
& J^{t+1}\left(S_{i}\right)=r_{i}+\gamma\left(\sum_{k} p_{i, k} J^{t}\left(s_{k}\right)\right)
\end{aligned}
$$

Remember, we have a converging function
We can stop when $\left|\mathrm{J}^{\mathrm{t}-1}\left(\mathrm{~s}_{\mathrm{i}}\right)-\mathrm{J}^{\mathrm{t}}\left(\mathrm{s}_{\mathrm{i}}\right)\right|_{\infty}<\varepsilon$

Infinity norm selects maximal element

## Example for $\gamma=0.9$

| $\mathrm{J}^{2}(\mathrm{Gr})=20+0.9^{*}\left(0.6^{*} 20+0.2^{*} 40\right.$ <br> $\left.+0.2^{*} 200\right)$ |  |  |  |  |
| :--- | :--- | :--- | :--- | :--- |
| t | $\mathrm{J}^{\mathrm{t}}(\mathrm{G})$ | $\mathrm{J}^{\mathrm{t}}(\mathrm{P})$ | $\mathrm{J}^{\mathrm{t}}(\mathrm{Goo})$ | $\mathrm{J}^{\mathrm{t}}(\mathrm{D})$ |
| 1 | 20 | 40 | 200 | 0 |
| 2 | $74^{\downarrow}$ | 87 | 362 | 0 |
| 3 | 141 | 135 | 493 | 0 |
| 4 | 209 | 182 | 600 | 0 |



## From MDPs to RL

- We still use the same Markov model with rewards and actions
- But there are a few differences:

1. We do not assume we know the Markov model
2. We adapt to new observations (online vs. offline)

- Examples:
- Game playing
- Robot interacting with enviroment
- Agents


## RL

- No actions
- With actions


## Scenario

- You wonder the world
- At each time point you see a state and a reward
- Your goal is to compute the sum of discounted rewards for each state



## Scenario

- You wonder the world
- At each time point you see a state and a reward
- Your goal is to compute the sum of discounted rewards for each state
- We will denote these by ${ }^{\text {jest }}\left(\mathrm{S}_{\mathrm{i}}\right)$



## Discounted rewards: $\gamma=0.9$

- Lets compute the discounted rewards for each time point: $\mathrm{t} 1: 4+0.9^{*} 0+0.9^{2 *} 2+0.9^{3 *} 2=7.1$
$\mathrm{t} 2: 0+0.9^{*} 2+0.9^{2 *} 2=3.4$
t3: $2+0.9 * 2$
$=3.8$
t4: $2+0$
$=2$
t5: 0
$=0$

| State | Observations | Mean |
| :--- | :--- | :--- |
| $S_{1}$ | 7.1 | 7.1 |
| $S_{2}$ | $3.4,2$ | 2.7 |
| $S_{3}$ | 3.8 | 3.8 |
| $S_{4}$ | 0 | 0 |



## Supervised learning for RL

- Type equation here.Observe set of states and rewards: (s(0),r(0)) ...(s(T),r(T))
- For $\mathrm{t}=0$... T compute discounted sum:

$$
J[t]=\sum_{i=t}^{T} \gamma^{i-t} r_{i}
$$

- Compute ${ }^{\text {Jest }}\left(\mathrm{s}_{\mathrm{i}}\right)=\left(\right.$ mean of $J(t)$ for $t$ such that $\left.s(t)=\mathrm{s}_{\mathrm{i}}\right)$

$$
\operatorname{Jest}\left[s_{i}\right]=\frac{\sum_{t \mid s t]=s_{i}}^{[1]}}{\# s[t]=s_{i}}
$$

We assume that we observe each state frequently enough and that we have many observations so that the final observations do not have a big impact on our prediction

## Algorithm for supervised learning

1. Initialize Counts $\left(\mathrm{s}_{\mathrm{i}}\right)=\mathrm{J}\left(\mathrm{s}_{\mathrm{i}}\right)=\operatorname{Disc}\left(\mathrm{s}_{\mathrm{i}}\right)=0$
2. Observe a state $s_{i}$ and a reward $r$
3. Counts $\left(\mathrm{s}_{\mathrm{i}}\right)=$ Counts $\left(\mathrm{s}_{\mathrm{i}}\right)+1$
4. $\operatorname{Disc}\left(\mathrm{s}_{\mathrm{i}}\right)=\operatorname{Disc}\left(\mathrm{s}_{\mathrm{i}}\right)+1$
5. For all states j

$$
\begin{aligned}
& \mathrm{J}\left(\mathrm{~s}_{\mathrm{j}}\right)=\mathrm{J}\left(\mathrm{~s}_{\mathrm{j}}\right)+\mathrm{r}^{*} \operatorname{Disc}\left(\mathrm{~s}_{\mathrm{j}}\right) \\
& \operatorname{Disc}\left(\mathrm{s}_{\mathrm{j}}\right)=\gamma^{*} \operatorname{Disc}\left(\mathrm{~s}_{\mathrm{j}}\right)
\end{aligned}
$$

6. Go to 2

At any time we can estimate $J^{*}$ by setting: Jest $\left(\mathrm{s}_{\mathrm{i}}\right)=\mathrm{J}\left(\mathrm{s}_{\mathrm{i}}\right) /$ Counts $\left(\mathrm{s}_{\mathrm{i}}\right)$

## Running time and space

- Each update takes $\mathrm{O}(\mathrm{n})$ where n is the number of states, since we are updating vectors containing entries for all states
- Space is also $\mathrm{O}(\mathrm{n})$

1. Convergences to true $\mathrm{J}^{*}$ can be proven
2. Can be more efficient by ignoring states for which

Disc() is very low already.

## Problems with supervised learning

- Takes a long time to converge
- Does not use all available data
- We can learn transition probabilities as well!


## Certainty-Equivalent (CE) Learning

- Lets try to learn the underlying Markov system's parameters



## CE learning

- We keep track of three vectors:

Counts(s): number of times we visited state s
$J(s)$ : sum of rewards from state s
Trans(i,j): number of time we transtiioned from state $\mathrm{s}_{\mathrm{i}}$ to state $\mathrm{s}_{\mathrm{j}}$

- When we visit state $s_{i}$, receive reward $r$ and move to state $s_{j}$ we do the following:

Counts $\left(\mathrm{s}_{\mathrm{i}}\right)=\operatorname{Counts}\left(\mathrm{s}_{\mathrm{i}}\right)+1$
$J\left(s_{i}\right)=J\left(s_{i}\right)+r$
$\operatorname{Trans}(\mathrm{i}, \mathrm{j})=\operatorname{Trans}(\mathrm{i}, \mathrm{j})+1$

## CE learning

- When we visit state $\mathrm{s}_{\mathrm{i}}$, receive reward r and move to state $\mathrm{s}_{\mathrm{j}}$ we do the following:
Counts $\left(\mathrm{s}_{\mathrm{i}}\right)=\operatorname{Counts}\left(\mathrm{s}_{\mathrm{i}}\right)+1$
$J\left(\mathrm{~s}_{\mathrm{i}}\right)=\mathrm{J}\left(\mathrm{s}_{\mathrm{i}}\right)+\mathrm{r}$
$\operatorname{Trans}(\mathrm{i}, \mathrm{j})=\operatorname{Trans}(\mathrm{i}, \mathrm{j})+1$

Using this we can estimate at any time the following parameters:

$$
\begin{aligned}
& \operatorname{Rest}\left(\mathrm{s}_{\mathrm{i}}\right)=\mathrm{J}\left(\mathrm{~s}_{\mathrm{i}}\right) / \operatorname{Counts}\left(\mathrm{s}_{\mathrm{i}}\right) \\
& \operatorname{Pest}(\mathrm{j} \mathrm{l})=\operatorname{Trans}(\mathrm{i}, \mathrm{j}) / \operatorname{Counts}\left(\mathrm{s}_{\mathrm{i}}\right)
\end{aligned}
$$

## Example: CE learning

Pest $(j \mid i)$

$$
\mathrm{R}^{\text {est }}\left(\mathrm{s}_{\mathrm{i}}\right)
$$

| State | Mean reward |
| :--- | :--- |
| $S_{1}$ | 4 |
| $S_{2}$ | 1 |
| $S_{3}$ | 2 |
| $S_{4}$ | 0 |


|  | $s 1$ | $s 2$ | $s 3$ | $s 4$ |
| :--- | :--- | :--- | :--- | :--- |
| s1 | 0 | 1 | 0 | 0 |
| s2 | 0 | 0 | 0.5 | 0.5 |
| s3 | 0 | 1 | 0 | 0 |
| s4 | 0 | 0 | 0 | 1 |



## CE learning

We can estimate at any time the following parameters:

$$
\begin{aligned}
& \left.\operatorname{Rest}^{\text {est }} \mathrm{s}_{\mathrm{i}}\right)=\mathrm{J}\left(\mathrm{~s}_{\mathrm{i}}\right) / \text { Counts }\left(\mathrm{s}_{\mathrm{i}}\right) \\
& \operatorname{Pest}\left(\mathrm{j}_{\mathrm{j}} \mathrm{i}\right)=\operatorname{Trans}(\mathrm{i}, \mathrm{j}) / \operatorname{Counts}\left(\mathrm{s}_{\mathrm{i}}\right)
\end{aligned}
$$

We now basically have an estimated which we can solve for all states $\mathrm{s}_{\mathrm{k}}$ :

$$
J^{e s t}\left(s_{k}\right)=r^{e s t}\left(s_{k}\right)+\gamma \sum_{j} p^{e s t}\left(s_{j} \mid s_{k}\right) J^{e s t}\left(s_{j}\right)
$$

## CE: Run time and space

## Run time

- Updates: O(1)
- Solving MDP:
- $O\left(n^{3}\right)$ using matrix inversion
- $\mathrm{O}\left(\mathrm{n}^{2 *} \#\right.$ it $)$ when using value iteration


## Space

- $\mathrm{O}\left(\mathrm{n}^{2}\right)$ for transition probabilities


## Improving CE: One backup

- We do the same updates and estimates as the original CE:

$$
\begin{aligned}
& \operatorname{Counts}\left(\mathrm{s}_{\mathrm{i}}\right)=\operatorname{Counts}\left(\mathrm{s}_{\mathrm{i}}\right)+1 \\
& \mathrm{~J}\left(\mathrm{~s}_{\mathrm{i}}\right)=\mathrm{J}\left(\mathrm{~s}_{\mathrm{i}}\right)+\mathrm{r} \\
& \operatorname{Trans}(\mathrm{i}, \mathrm{j})=\operatorname{Trans}(\mathrm{i}, \mathrm{j})+1
\end{aligned}
$$

$$
\begin{aligned}
& R^{\text {est }}\left(\mathrm{s}_{\mathrm{i}}\right)=\mathrm{J}\left(\mathrm{~s}_{\mathrm{i}}\right) / \text { Counts }\left(\mathrm{s}_{\mathrm{i}}\right) \\
& \text { Pest } \left.^{\mathrm{j}} \mathrm{j} \mathrm{i}\right)=\operatorname{Trans}(\mathrm{i}, \mathrm{j}) / \operatorname{Counts}\left(\mathrm{s}_{\mathrm{i}}\right)
\end{aligned}
$$

- But we do not carry out the full value iteration
- Instead, we only update ${ }^{\text {Jest }}\left(\mathrm{s}_{\mathrm{i}}\right)$ for the current state:

$$
J^{e s t}\left(s_{i}\right)=r^{\text {est }}\left(s_{i}\right)+\gamma \sum_{j} p^{e s t}\left(s_{j} \mid s_{i}\right) J^{e s t}\left(s_{j}\right)
$$

## CE one backup: Run time and space

## Run time

- Updates: O(1)
- Solving MDP:
- O(1) just update current state

Space

- $\mathrm{O}\left(\mathrm{n}^{2}\right)$ for transition probabilities
- Still a lot of memory, but much more efficient
- Can prove convergence to optimal solution (but slower than CE)


## Summary so far

- Three methods

| Method | Time | Space |
| :--- | :--- | :--- |
| Supervised learning | $\mathrm{O}(\mathrm{n})$ | $\mathrm{O}(\mathrm{n})$ |
| CE learning | $\mathrm{O}\left(\mathrm{n}^{2 \star} \# \mathrm{it}\right)$ | $\mathrm{O}\left(\mathrm{n}^{2}\right)$ |
| One backup CE | $\mathrm{O}(1)$ | $\mathrm{O}\left(\mathrm{n}^{2}\right)$ |

## Temporal difference (TD) learning

- Goal: Same efficiency as one backup CE while much less space
- We only maintain the Jest array.
- Assume we have Jest $\left(s_{1}\right)$... Jest $\left(s_{n}\right)$. If we observe a transition from state $\mathrm{s}_{\mathrm{i}}$ to state $\mathrm{s}_{\mathrm{j}}$ and a reward r , we update using the following rule:

$$
J^{e s t}\left(s_{i}\right)=(1-\alpha) J^{e s t}\left(s_{i}\right)+\alpha\left(r+\gamma j^{e s t}\left(s_{j}\right)\right)
$$

## Temporal difference (TD) learning

- Assume we have ${ }^{\text {Jest }}\left(\mathrm{s}_{1}\right)$... Jest $\left(\mathrm{s}_{\mathrm{n}}\right)$. If we observe a transition from state $\mathrm{s}_{\mathrm{i}}$ to state $\mathrm{s}_{\mathrm{j}}$ and a reward r , we update using the following rule:

$$
J^{\text {est }}\left(s_{i}\right)=(1-\alpha) J^{e s t}\left(s_{i}\right)+\alpha\left(r+\gamma j^{e s t}\left(s_{j}\right)\right)
$$

parameter to determine how much
weight we place on current
observation
We have seen similar update rule before, as always, choosing $\alpha$ is an issue

## Convergence

- TD learning is guaranteed to converge if:
- All states are visited often
- And: $\sum_{t} \alpha_{t}=\infty$

$$
\sum_{t} \alpha_{t}^{2}<\infty
$$

For example, $\alpha_{\mathrm{t}}=\mathrm{C} / \mathrm{t}$ for some constant C would satisfy both requirements

## TD: Complexity and space

- Time to update: O(1)
- Space: O(n)

| Method | Time | Space |
| :--- | :--- | :--- |
| Supervised <br> learning | $\mathrm{O}(\mathrm{n})$ | $\mathrm{O}(\mathrm{n})$ |
| CE learning | $\mathrm{O}\left(\mathrm{n}^{2 \star} \#\right.$ it $)$ | $\mathrm{O}\left(\mathrm{n}^{2}\right)$ |
| One backup CE | $\mathrm{O}(1)$ | $\mathrm{O}\left(\mathrm{n}^{2}\right)$ |

## RL

- No actions $\sqrt{ }$
- With actions


## Policy learning

- So far we assumed that we cannot impact the outcome transition.
- In real world situations we often have a choice of actions we take (as we discussed for MDPs).
- How can we learn the best policy for such cases?



## Policy learning using CE : Example

| $\mathrm{R}^{\text {est }}\left(\mathrm{s}_{\mathrm{i}}\right)$ |  | Pest(jili,a) |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  | s1 | s2 | s3 | s4 |
| State | Mean reward | s1,A | 0 | 0 | 1 | 0 |
| $\mathrm{S}_{1}$ | 4 | s1,B | 0 | 1 | 0 | 0 |
| $\mathrm{S}_{2}$ | 4/3 | s2 | 0 | 0 | 2/3 | 1/3 |
| $\mathrm{S}_{3}$ | 2.5 | s3 | 1/3 | 1/3 | 0 | 1/3 |
| $\mathrm{S}_{4}$ | 0 | s4 | 0 | 1 | 0 | 0 |



## Policy learning using CE

We can easily update CE by setting:

$$
J^{e s t}\left(s_{k}\right)=r^{e s t}\left(s_{k}\right)+\max _{a}\left[\gamma \sum_{j} p^{e s t}\left(s_{j} \mid s_{k}, a\right) J^{e s t}\left(s_{j}\right)\right]
$$

We revise our transition model to include actions

## Policy learning for TD

- TD is model free
- We can adjust TD to learn policies by defining the $Q$ function:
- $\mathrm{Q}^{*}\left(\mathrm{~s}_{\mathrm{i}}, \mathrm{a}\right)=$ expected sum of future (discounted) rewards if we start at state $\mathrm{s}_{\mathrm{i}}$ and take action a
- When we take a specific action a in state $\mathrm{s}_{\mathrm{i}}$ and then transition to state $\mathrm{s}_{\mathrm{j}}$ we can update the Q function directly by setting:
$Q^{e s t}\left(S_{i}, a\right)=(1-\alpha) Q^{e s t}\left(S_{i}, a\right)+\alpha\left(r_{i}+\gamma \max _{a^{\prime}} Q^{e s t}\left(S_{j}, a^{\prime}\right)\right)$

Instead of the Jest vector we maintain the Qest matrix, which is a rather sparse $n$ by matrix ( n states and $m$ actions)

## Choosing the next action

- We can select the action that results in the highest expected sum of future rewards
- But that may not be the best action. Remember, we are only sampling from the distribution of possible outcomes. We do not want to avoid potentially beneficial actions.
- Instead, we can take a more probabilistic approach:

The probability we will use action a

## Choosing the next action

- Instead, we can take a more probabilistic approach:

$$
p(a) \propto \exp \left(-\frac{Q^{e s t}\left(s_{i}, a\right)}{f(t)}\right)
$$

- We can initialize $Q$ values to be high to increase the likelihood that we will explore more options
- It can be shown that $Q$ learning converges to optimal policy


## What you should know

- Strategies for computing with expected rewards
- Strategies for computing rewards and actions
- Q learning

