## SMU: Lecture 4

Monday, March 7, 2022
(Heavily inspired by the Stanford RL Course of Prof. Emma Brunskill, but all potential errors are mine.)

## Plan for Today

- Recap of important concepts from lectures 1, 2 and 3.
- Value function approximation.
- Control with value function approximation.


## Part 1: Where are we? (Recap from the previous lectures)

## Markov Decision Process

- Markov decision process = Markov reward process + Actions
- An MDP is given by:
- A set of states $S$.
- A set of actions $A$.

$$
\text { A transition model } P\left[X_{t+1}=s^{\prime} \mid X_{t}=s, A_{t}=a\right] \underbrace{=P\left(s^{\prime} \mid s, a\right)}_{\text {notation }}
$$

- A reward $R(s, a)=\mathbb{E}\left[R_{t} \mid X_{t}=s, A_{t}=a\right]$, i.e. the expected reward that the agent receives when performing action $a$ in state $s$.
- Discount factor $\gamma$.


## Policy

- Policy determines which action to take in each state $s$.
- It can be either deterministic or random - that is also why policy will not simply be a function from states to actions.
- We define policy: $\pi(a \mid s)=P\left(A_{t}=a \mid X_{t}=s\right)$.
- Example (policy for our ant ${ }_{\text {pep }}$ ):
- $A=\{$ left, right $\}$
- $\pi($ left $\mid 1)=0, \pi($ right $\mid 1)=1, \pi($ left $\mid 2)=0.5, \pi($ right $\mid 1)=0.5, \ldots$


## State Value Function of MDP

## General case:

$$
V^{\pi}(s)=\sum_{a \in A} \pi(a, s) \cdot\left[R(s, a)+\gamma \cdot \sum_{s^{\prime} \in S} P\left(s^{\prime} \mid s, a\right) \cdot V^{\pi}\left(s^{\prime}\right)\right]
$$

Version for deterministic policy:

$$
V^{\pi}(s)=R(s, \pi(s))+\gamma \cdot \sum_{s^{\prime} \in S} P\left(s^{\prime} \mid s, \pi(s)\right) \cdot V^{\pi}\left(s^{\prime}\right)
$$

## MDP Control Problem

How to find $\pi^{*}(s)=\arg \max V^{\pi}(s)$ ???

## State-Action Value Q

- Definition:

$$
Q^{\pi}(s, a)=R(s, a)+\gamma \cdot \sum_{s^{\prime} \in S} P\left(s^{\prime} \mid s, a\right) \cdot V^{\pi}\left(s^{\prime}\right)
$$

- Intuition:
- The value of the return that we obtain if we first take the action $a$ in the state $s$ and then follow the policy $\pi$ (including when we visit $s$ again).
- Think of it as perturbing the policy $\pi$ - we deviate from following the policy $\pi$ only in the first step in $s$.


## Policy Improvement Step

- Given: An MDP and a policy $\pi_{i}$ that we want to improve (if possible).
- DO:
- For all $s \in S$, compute $Q^{\pi_{i}}(s, a)$ as defined on the previous slide, i.e. $Q^{\pi_{i}}(s, a)=R(s, a)+\gamma \cdot \sum_{s^{\prime} \in S} P\left(s^{\prime} \mid s, a\right) \cdot V^{\pi_{i}}\left(s^{\prime}\right)$.
- Compute new policy for all $s \in S$ : Here, we use the fact that our policy is deterministic for simpler notation (treating policy as a function). Using our previous notation we could write:

$$
\pi_{i+1}(s)=\arg \max _{a \in S} Q^{\pi_{i}(s, a)}
$$

$$
\pi(a \mid s)= \begin{cases}1 & \text { if } a=\arg \max _{a \in A} Q^{\pi_{i}}(s, a) \\ 0 & \text { otherwise }\end{cases}
$$

## Example

Agent:


## Rewards??



Actions are given:

$$
A=\{l, r\}
$$

Policy is given, e.g.:

$$
\begin{aligned}
& \pi(l \mid a)=0.2, \pi(r \mid a)=0.8 \\
& \pi(l \mid b)=0.3, \pi(r \mid b)=0.7
\end{aligned}
$$

## Problem: Model-Free Policy Evaluation

- Given a policy and an MDP with unknown parameters (or generally an environment with which we can interact), estimate the value function.


## Model-Free Control

- Given a policy and an MDP with unknown parameters (or generally an environment with which we can interact), find the optimal policy $\pi$.


## Three Methods Last Time

- Monte Carlo Control, SARSA and Q-Learning.
- All three using the concept of $\varepsilon$-greedy policy.


## MC On Policy Improvement

Initialize: $G(s, a)=0, N(s, a)=0, Q(s, a)=0$ for all $s \in S, a \in A$.
Initialize: $\varepsilon=1, k=1$
For $i=1, \ldots, N$ :
Sample episode $e_{i}:=s_{i, 1}, a_{i, 1}, r_{i, 1}, s_{i, 2}, a_{i, 2}, r_{i, 2}, \ldots, s_{i, T_{i}}$ given $\pi_{k}$.
For each time step $1 \leq t \leq T_{i}$ :
(If $t$ is the first occurrence of state $s$ in the episode $e_{i}$ - Use this if you want first-visit MC)
$s_{t}$ is the state visited at time $t$ in the episode $e_{i}$
$a_{t}$ is the action taken at time $t$ in the episode $e_{i}$
$g_{i, t}:=r_{i, t}+\gamma \cdot r_{i, t+1}+\gamma^{2} \cdot r_{i, t+2}+\ldots+\gamma^{T_{i}-t} \cdot r_{i, T_{i}}$
$N(s):=N(s)+1 /{ }^{*}$ Increment total visits counter */
$G\left(s_{t}, a_{t}\right):=G\left(s_{t}, a_{t}\right)+g_{i, 1} / *$ Increment total return counter */
$Q\left(s_{t}, a_{t}\right):=G\left(s_{t}, a_{t}\right) / N\left(s_{t}, a_{t}\right) / *$ Update current estimate *

## EndFor

$k=k+1, \varepsilon=1 / k$
$\pi_{k}=\varepsilon-$ greedy policy w.r.t. $Q$

## SARSA

- SARSA is an on-policy algorithm.

1. Initialize: set $\pi$ to be some $\varepsilon$-greedy policy, set $t=0$
2. Sample $a$ using the distribution given by $\pi_{0}$ in the state $s_{0}$ (for sampling, we will use the notation $a \sim \pi(s)$ ). Take the action $a$ and observe $r_{0}, s_{1}$.
3. While $s_{t}$ is not a terminal state:
4. Take action $a \sim \pi\left(s_{t}\right)$ and observe $r_{t+1}, s_{t+a}$.
5. $Q\left(s_{t}, a_{t}\right):=Q\left(s_{t}, a_{t}\right)+\alpha\left(r_{t}+\gamma Q\left(s_{t+1}, a_{t+1}\right)-Q\left(s_{t}, a_{t}\right)\right)$
6. $\pi:=\varepsilon$-greedy $(Q)$
7. Set $t:=t+1$. Update $\varepsilon, \alpha / *$ see next slides */

## Q-Learning

- Q-Learning is an off-policy algorithm.

1. Initialize: set $\pi$ to be some $\varepsilon$-greedy policy, set $t=0$
2. Sample $a$ using the distribution given by $\pi_{0}$ in the state $s_{0}$ (for sampling, we will use the notation $a \sim \pi(s)$ ). Take the action $a$ and observe $r_{0}, s_{1}$.
3. While $s_{t}$ is not a terminal state:
4. Take action $a \sim \pi\left(s_{t}\right)$ and observe $r_{t+1}, s_{t+1}$.
5. $Q\left(s_{t}, a_{t}\right):=Q\left(s_{t}, a_{t}\right)+\alpha\left(r_{t}+\gamma \max _{a \in A} Q\left(s_{t+1}, a\right)-Q\left(s_{t}, a_{t}\right)\right)$
6. $\pi:=\varepsilon$-greedy $(Q)$
7. Set $t:=t+1$. Update $\varepsilon, \alpha / *$ see next slides */

## $\varepsilon$-Greedy Policy

$$
\pi(a \mid s)= \begin{cases}1-\varepsilon+\frac{\varepsilon}{|A|} & \text { when } a=\arg \max _{a \in A}^{\text {We assume ties are d }} Q(s, a) \\ \frac{\varepsilon}{|A|} & \text { when } a \neq \arg \max _{a \in A} Q(s, a)\end{cases}
$$

## Part 2: RL with Function <br> Approximation (Problem Statement)

## Limitations of What We Saw So Far

- In the previous lectures, we assumed discrete MDPs with number of states that was not too large (i.e. the set $S$ was not too large).
- Now imagine that we want to learn to play Atari games (which is what DeepMind did!) and we want to do it form the pixel inputs. How many states would we need if we wanted to use what we learned in the previous lectures? ... Then we would need at least $128^{160 \cdot 192}$ states ( 128 colors with resolution $160 \times 192$ pixels).
- What we need is function approximation.


## Besic ce?

- Do not represent the state value function $V$ or the state-action value function $Q$ explicitly.
- Represent the state value function $V(s)$ or the state-action value function $Q(s, a)$ approximately using a function from some parametrized family, e.g. as a neural network, linear function, decision tree...



$w \quad \rightarrow \hat{Q}(s, a ; w)$
$a, a \in\{$ left, right, up, down $\} \rightarrow$


## Part 3: Some Background

## Gradient Descent (1/3)

- A method for finding a (local) optimum of a function.
- In our setting, we want to find $\mathbf{w} \in \mathbb{R}^{d}$ that is a local minimum of a function $J(\mathbf{w})$.
- We do that using gradient descent.


## Gradient Descent (2/3)

Gradient: $\quad \nabla J(\mathbf{w})=\left(\frac{\partial J}{\partial w_{1}}(\mathbf{w}), \frac{\partial J}{\partial w_{2}}(\mathbf{w}), \ldots, \frac{\partial J}{\partial w_{d}}(\mathbf{w})\right)$

## Example:

$$
J(\mathbf{w})=w_{1} \cdot w_{2}+w_{1}, \mathbf{w} \in \mathbb{R}^{2} .
$$

Then
$\nabla J(\mathbf{w})=\left(w_{2}+1, w_{1}\right)$.

## Gradient Descent (3/3)

Gradient descent update rule:

$$
\mathbf{w}_{n+1}=\mathbf{w}_{n}-\alpha \cdot \nabla J\left(\mathbf{w}_{n}\right)
$$

(gradient descent algorithm iterates this rule).

## A Useful Property of Mean Squared Loss

Let $Y_{1}, Y_{2}, \ldots, Y_{n}$ be an independent sample from some distribution with expected value $\mu=\mathbb{E}\left[Y_{i}\right], \forall i$.

What is the value $y$ ( $\sim$ prediction) that minimizes the mean squared error
$\frac{1}{n} \sum_{i=1}^{n}\left(Y_{i}-y\right)^{2}$ ?
It is the sample average $y=\frac{1}{n} \sum_{i=1}^{n} Y_{i}$, which, for $n \rightarrow \infty$, converges to the mean $\mu$.
Consequence: Learning a predictor under mean squared loss leads to learning a predictor for conditional expectation (we will explain later what it means for RL).

## Stochastic Gradient Descent

- We want to optimize a function $J(\mathbf{w})$ of the form $J(\mathbf{w})=\mathbb{E}[g(X ; \mathbf{w})]$ where $X$ is a random variable.
- We assume that we can sample from the distribution w.r.t. which the expectation is taken.
- Stochastic gradient descent uses samples to approximate the gradient of $J(\mathbf{w})$ using just one sample (SGD can also use a mini-batch of multiple samples but we will not consider it now for simplicity) and estimates the gradient of $J$ as:

$$
\begin{gathered}
\nabla J(\mathbf{w}) \approx \nabla g(X ; \mathbf{w}) \\
\text { (instead of } \nabla \mathbb{E}[g(X ; \mathbf{w})]) .
\end{gathered}
$$

- Assuming that we can exchange the order of expectation and taking gradients (which we can when $g$ is well-behaved), the expected SGD step is the same as the full gradient of J.


## Warm-Up: Learning to "Compress" $V^{\pi}(s),(1 / 3)$

- Suppose that we know $V^{\pi}(s)$ and can query it but yet want to learn an approximation of it... using a parametric function $\hat{V}^{\pi}(s ; \mathbf{w}) \ldots$
- We will use mean-squared error to measure how good the approximation is, i.e.:

$$
J(\mathbf{w})=\mathbb{E}_{\pi}\left[\left(V^{\pi}(X)-\hat{V}^{\pi}(X ; \mathbf{w})\right)^{2}\right] .
$$

- How could we train the approximation using SGD?


## Warm-Up: Learning to "Compress" $V^{\pi}(s),(2 / 3)$

- We will use mean-squared error to measure how good the approximation is, i.e.:

$$
J(\mathbf{w})=\mathbb{E}_{\pi}\left[\left(V^{\pi}(X)-\hat{V}^{\pi}(X ; \mathbf{w})\right)^{2}\right]
$$

- How could we train the approximation using SGD (using 1 example in each minibatch)?


## While (some stopping condition):

Sample a state $s$ and compute the gradient of $\hat{J}_{s}(\mathbf{w})=\left(V^{\pi}(s)-V(s ; \mathbf{w})\right)^{2}$ :

$$
\nabla \hat{J}_{s}(\mathbf{w})=-2\left(V^{\pi}(s)-V^{\pi}(s ; \mathbf{w})\right) \cdot \nabla V^{\pi}(s ; \mathbf{w})=2\left(V^{\pi}(s ; \mathbf{w})-V^{\pi}(s)\right) \cdot \nabla V^{\pi}(s ; \mathbf{w})
$$

Take the gradient step:

$$
\mathbf{w}:=\mathbf{w}-\alpha \cdot 2\left(V^{\pi}(s ; \mathbf{w})-V^{\pi}(s)\right) \cdot \nabla V(s ; \mathbf{w})
$$

## Warm-Up: Learning to "Compress" $V^{\pi}(s),(3 / 3)$

- But in reality we will not have access to $V^{\pi}(s)$ !
- So we cannot compute the gradient step: $\mathbf{w}:=\mathbf{w}-\alpha \cdot 2\left(V^{\pi}(s ; \mathbf{w})-V^{\pi}(s)\right) \cdot \nabla V(s ; \mathbf{w}) \ldots$
- We will therefore need to combine SGD with what we saw in the previous lectures...


## Part 4: Let’s Make It Concrete (Some Parametric Families)

## State Representation

- States will be represented by feature vectors.
- The feature vector of a state $s$ will be denoted as $\mathbf{x}(s)$ and we can think of it as a function mapping states to some vector space, e.g. $\mathbb{R}^{d}$, i.e.

$$
\mathbf{x}(s)=\left(x_{1}(s), x_{2}(s), \ldots, x_{d}(s)\right)^{T}
$$

- Examples:
- Atari: the feature vector can, e.g., contain the intensities of the pixels (concatenated).
- Pole balancing: physical features such as velocities, angles...


## Linear Functions

- Scalar product of a weight vector with the feature vector, which represents the state:

$$
\hat{V}^{\pi}(s ; \mathbf{w})=\mathbf{w}^{T} \mathbf{x}(s) .
$$

- Linear function approximations can work well but need good features (which requires feature engineering).


## Neural Networks

- Neural network (well, you know them):

- In this lecture we will think of neural networks simply as blackboxes $g(\mathbf{x} ; \mathbf{w})$ which we can evaluate and for which we can compute the gradients $\nabla_{\mathbf{w}} g(\mathbf{x} ; \mathbf{w})$ efficiently (we will usually omit the subscript $\mathbf{w}$ from $\nabla_{\mathbf{w}}$ when it is clear from the context).
- In particular, the approximation will have the form $V^{\pi}(s ; \mathbf{w})=g(\mathbf{x}(s) ; \mathbf{w})$, where $g$ is some neural network...


## Part 5: Policy Evaluation with Function Approximation

## Monte-Carlo Value Function Approximation

We know from the previous lectures that the return
$G_{t}^{\pi}=R\left(X_{t}, \pi\left(X_{t}\right)\right)+\gamma \cdot R\left(X_{t+1}, \pi\left(X_{t+1}\right)\right)+\ldots$ is an unbiased estimate (but with high variance) of $V^{\pi}\left(X_{t}\right) \ldots$ Conditioning on the value of $X_{t}$, we then arrive at the Monte-Carlo estimators, i.e. we have $V^{\pi}(s)=\mathbb{E}\left[G_{t} \mid X_{t}=s\right]$.

Basic Idea (not yet complete... wait for the next slide): We can frame the value function approximation problem as a supervised learning problem under MSE loss:
Sample an episode under policy $\pi$ : $s_{1}, a_{1}, r_{1}, s_{2}, a_{2}, r_{2}, \ldots, s_{T}$
Training examples: $\left[s_{1}, g_{1}\right],\left[s_{2}, g_{2}\right], \ldots,\left[s_{T-1}, g_{T-1}\right]$, where $g_{i}$ denotes the return from the episode from time $i$.

## First/Every-Visit Monte-Carlo Value Function Approximation

Initialize: w = some initialization....
For $i=1, \ldots, N$ :
Sample episode $e_{i}:=s_{i, 1}, a_{i, 1}, r_{i, 1}, s_{i, 2}, a_{i, 2}, r_{i, 2}, \ldots, s_{i, T_{i}}$.
For each time step $1 \leq t \leq T_{i}$ :
If $t$ is the first occurrence of state $s$ in the episode $e_{i} / *$ This is for first-visit MC */
$s$ is the state visited at time $t$ in the episode $e_{i}$

$$
\begin{aligned}
& g_{i, t}:=r_{i, t}+\gamma \cdot r_{i, t+1}+\gamma^{2} \cdot r_{i, t+2}+\ldots+\gamma^{T_{i}-t} \cdot r_{i, T_{i}} \\
& / * \text { SGD step */ } \\
& \mathbf{w}:=\mathbf{w}-\alpha \cdot\left(V^{\pi}\left(\mathbf{x}\left(s_{t}\right) ; \mathbf{w}\right)-g_{t}\right) \cdot \nabla V\left(\mathbf{x}\left(s_{t}\right) ; \mathbf{w}\right)
\end{aligned}
$$

## MC Value Function Evaluation with Linear Function Approximation

We have

$$
\begin{aligned}
& \hat{V}^{\pi}(s ; \mathbf{w})=\mathbf{w}^{T} \mathbf{x}(s) \\
& \nabla \hat{V}^{\pi}(s ; \mathbf{w})=\mathbf{x}(s) .
\end{aligned}
$$

Initialize: w = some initialization....
For $i=1, \ldots, N$ :
Sample episode $e_{i}:=s_{i, 1}, a_{i, 1}, r_{i, 1}, s_{i, 2}, a_{i, 2}, r_{i, 2}, \ldots, s_{i, T_{i}}$
For each time step $1 \leq t \leq T_{i}$ :
If $t$ is the first occurrence of state $s$ in the episode $e_{i} /{ }^{*}$ This is for first-visit MC */
$s$ is the state visited at time $t$ in the episode $e_{i}$

$$
g_{i, t}:=r_{i, t}+\gamma \cdot r_{i, t+1}+\gamma^{2} \cdot r_{i, t+2}+\ldots+\gamma^{T_{i}-t} \cdot r_{i, T_{i}}
$$

/* SGD step */

$$
\mathbf{w}:=\mathbf{w}-\alpha \cdot\left(\mathbf{w}^{T} \mathbf{x}\left(s_{t}\right)-g_{t}\right) \cdot \mathbf{x}\left(s_{t}\right)
$$

## Intuition About Why It Works

- Recall that what we want to estimate is $V^{\pi}(s)=\mathbb{E}\left[G_{t} \mid X_{t}=s\right]$, i.e. expected value.
- When using first-visit MC, each of the training examples $\left[s_{t}, g_{t}\right]$ is an unbiased (but very noisy!) estimate of $V^{\pi}(s)$. But when we use these examples and try to find a best mean-squared-error fit then we are estimating their expectation which equals $V^{\pi}(s)$. And that is why it works...


## Convergence of MC VFA (1/3)

- Definition (On-Policy Distribution): Given an MDP and a policy $\pi$, we define on-policy distribution $P_{o n p}^{\pi}$ as follows.
- In non-episodic settings: $P_{o n p}^{\pi}$ is the stationary distribution of the MRP that is given by the MDP and the policy (recall MDP + policy $=M R P$ ).
- In episodic settings: $P_{o n p}^{\pi}$ depends also on the distribution of the initial states $P_{\text {init }}$ (see Sutton's book for details).
- In what follows, we denote the on-policy distribution by $P_{o n p}^{\pi}$.


## Convergence of MC VFA (2/3)

- Definition: Mean squared error of value function approximation is defined as

$$
\operatorname{MSVE}_{\pi}(\mathbf{w})=\sum_{s \in S} P_{o n p}^{\pi}(s) \cdot\left(V^{\pi}(s)-\hat{V}^{\pi}(s ; \mathbf{w})\right)^{2}
$$

which is the same as

$$
\operatorname{MSVE}_{\pi}(\mathbf{w})=\mathbb{E}_{X \sim P_{\text {onp }}^{\pi}}\left[\left(V^{\pi}(s)-\hat{V}^{\pi}(s ; \mathbf{w})\right)^{2}\right]
$$

## Convergence of MC VFA (3/3)

- Theorem: Assume that $\hat{V}^{\pi}(s ; \mathbf{w})=\mathbf{w}^{T} \mathbf{x}(s)$, i.e. we are assuming linear function approximation. Then MC VFA converges to weights that are optimal in the sense that they minimize $\operatorname{MSVE}_{\pi}(\mathbf{w})$.
- Caution: This theorem holds for linear function approximation, not for general functions! We do not have such guarantees for, e.g., arbitrary neural networks.


## Temporal Difference VFA (1/5)

- For temporal difference learning, we had the following update rule:

$$
V^{\pi}\left(s_{t}\right):=V^{\pi}\left(s_{t}\right)+\alpha \cdot(\underbrace{r_{t}+\gamma \cdot V^{\pi}\left(s_{t+1}\right)}_{\text {TD-target }}-V^{\pi}\left(s_{t}\right))
$$

- Now, we will want to have a similar update rule but for the case where $V^{\pi}(s)$ is only approximated by $V^{\pi}(s ; \mathbf{w})$.


## Temporal Difference VFA (2/5)

Recall the Bellman equation (for simplicity, we are showing it for deterministic policy): $V^{\pi}(s)=R(s, \pi(s))+\gamma \cdot \sum_{s^{\prime} \in S} P\left(s^{\prime} \mid s, \pi(s)\right) \cdot V\left(s^{\prime}\right)$
which is the same as:

$$
V^{\pi}(s)=R(s, \pi(s))+\gamma \cdot \mathbb{E}\left[V^{\pi}\left(X_{t+1}\right) \mid X_{t}=s\right] .
$$

We can turn the system of equations above into the following minimization problem:

$$
\min _{\mathbf{V}^{\pi}} \sum_{s \in S} P_{s t a t}(s) \cdot \mathbb{E}\left[\left(R(s, \pi(s))+\gamma \cdot V^{\pi}\left(X_{t+1}\right)-V^{\pi}(s)\right)^{2} \mid X_{t}=s\right]
$$

## Temporal Difference VFA (3/5)

Next we replace $V^{\pi}(s)$ by its approximation $\hat{V}^{\pi}(s ; \mathbf{w})$, yielding:

$$
\min _{\mathbf{V}^{\pi}} \sum_{s \in S} P_{s t a t}(s) \cdot \mathbb{E}\left[\left(R(s, \pi(s))+\gamma \cdot \hat{V}^{\pi}\left(X_{t+1} ; \mathbf{w}\right)-\hat{V}^{\pi}(s ; \mathbf{w})\right)^{2} \mid X_{t}=s\right]
$$

Now, instead of the stationary distribution, we will just take the distribution of states as they come in an episode and instead of the expectation we will use the tuple ( $s_{t}, a_{t}, r_{t}, s_{t+1}$ ) we get in the current episode (as is common in TDlearning). That will lead us to the minimization problem:

$$
\min _{\mathbf{w}}\left(R\left(s_{t}, r_{t}\right)+\gamma \cdot \hat{V}^{\pi}\left(s_{t+1} ; \mathbf{w}\right)-\hat{V}^{\pi}\left(s_{t} ; \mathbf{w}\right)\right)^{2}
$$

## Temporal Difference VFA (4/5)

We need to solve:

$$
\begin{aligned}
& \min _{\mathbf{w}}\left(R\left(s_{t}, r_{t}\right)+\gamma \cdot \hat{V}^{\pi}\left(s_{t+1} ; \mathbf{w}\right)-\hat{V}^{\pi}\left(s_{t} ; \mathbf{w}\right)\right)^{2} \text {. Denoting } \\
& J(\mathbf{w})=\left(R\left(s_{t}, r_{t}\right)+\gamma \cdot \hat{V}^{\pi}\left(s_{t+1} ; \mathbf{w}\right)-\hat{V}^{\pi}\left(s_{t} ; \mathbf{w}\right)\right)^{2}
\end{aligned}
$$

we have

$$
\nabla J(\mathbf{w})=2\left(R\left(s_{t}, r_{t}\right)+\gamma \cdot \hat{V}^{\pi}\left(s_{t+1} ; \mathbf{w}\right)-\hat{V}^{\pi}\left(s_{t} ; \mathbf{w}\right)\right) \cdot\left(\gamma \cdot \nabla \hat{V}^{\pi}\left(s_{t+1} ; \mathbf{w}\right)-\nabla \hat{V}^{\pi}\left(s_{t} ; \mathbf{w}\right)\right)
$$

But this is not what TD with function approximation does! TD VFA is a socalled semigradient method. It does not consider the contribution of $\nabla \hat{V}^{\pi}\left(s_{t+1} ; \mathbf{w}\right)$ and considers it fixed.

## Temporal Difference VFA (4/5)

We need to solve:

$$
\begin{aligned}
& \min _{\mathbf{w}}\left(R\left(s_{t}, r_{t}\right)+\gamma \cdot \hat{V}^{\pi}\left(s_{t+1} ; \mathbf{w}\right)-\hat{V}^{\pi}\left(s_{t} ; \mathbf{w}\right)\right)^{2} \text {. Denoting } \\
& J(\mathbf{w})=\left(R\left(s_{t}, r_{t}\right)+\gamma \cdot \hat{V}^{\pi}\left(s_{t+1} ; \mathbf{w}\right)-\hat{V}^{\pi}\left(s_{t} ; \mathbf{w}\right)\right)^{2}
\end{aligned}
$$

we have

$$
\nabla J(\mathbf{w})=2\left(R\left(s_{t}, r_{t}\right)+\gamma \cdot \hat{V}^{\pi}\left(s_{t+1} ; \mathbf{w}\right)-\hat{V}^{\pi}\left(s_{i} ; \mathbf{w}\right)\right) \cdot\left(\gamma \cdot \nabla \hat{V}^{\pi}\left(s_{t+1} ; \mathbf{w}\right)-\nabla \hat{V}^{\pi}\left(s_{i} ; \mathbf{w}\right)\right)
$$

But this is not what TD with function approximation does! TD VFA is a socalled semigradient method. It does not consider the contribution of $\nabla \hat{V}^{\pi}\left(s_{t+1} ; \mathbf{w}\right)$ and considers it fixed.

## Temporal Difference VFA (5/5)

The TD update rule for value function approximation is:

$$
\mathbf{w}:=\mathbf{w}+\alpha\left(r_{t}+\gamma \cdot \hat{V}^{\pi}\left(s_{t+1} ; \mathbf{w}\right)-\hat{V}^{\pi}\left(s_{t} ; \mathbf{w}\right)\right) \cdot \nabla \hat{V}^{\pi}\left(s_{t} ; \mathbf{w}\right)
$$

## TD VFA with Linear Functions

We have

$$
\begin{aligned}
& \hat{V}^{\pi}(s ; \mathbf{w})=\mathbf{w}^{T} \mathbf{x}(s) \\
& \nabla \hat{V}^{\pi}(s ; \mathbf{w})=\mathbf{x}(s)
\end{aligned}
$$

The TD update rule for value function approximation with linear function is:

$$
\mathbf{w}:=\mathbf{w}+\alpha\left(r_{t}+\gamma \cdot \mathbf{w}^{T} \mathbf{x}\left(s_{t+1}\right)-w^{T} \mathbf{x}\left(s_{t}\right)\right) \mathbf{x}\left(s_{t}\right) .
$$

## Convergence of TD VFA with Linear Functions

- As for MC VFA, we will use the on-policy distribution $P_{o n p}^{\pi}$ and define the mean squared error w.r.t. it, that is...

$$
\operatorname{MSVE}_{\pi}(\mathbf{w})=\sum_{s \in S} P_{o n p}^{\pi}(s) \cdot\left(V^{\pi}(s)-\hat{V}^{\pi}(s ; \mathbf{w})\right)^{2}
$$

- Theorem: Let $\mathbf{w}_{T D}$ be the weight vector to which TD VFA converges. Then it holds:

$$
\operatorname{MSV}_{\pi}\left(\mathbf{w}_{T D}\right) \leq \frac{1}{1-\gamma} \cdot \min _{\mathbf{w}} \operatorname{MSVE}_{\pi}(\mathbf{w}) .
$$

- Recall that for MC VFA with linear functions we had convergence of mean squared error to $\min _{\mathbf{w}} \operatorname{MSVE}_{\pi}(\mathbf{w})$.
w


## Part 6: Control with Function Approximation

## Basic Idea

- Same ideas, just plugging them into what we were doing in the last lecture, but there are caveats...
- Instead of approximating $V^{\pi}$, we need to approximate $Q^{\pi}(s, a)$ from which we can then extract the best action to take in a given state $s$.
- The algorithms are similar to those we saw last week (MC, SARSA, QLearning). Important: the idea of using $\varepsilon$-greedy policies. The motivation is the same but we use $Q^{\pi}(s, a ; \mathbf{w})$.


## Representing State-Action Pairs

- For control RL problems, we need to encode both states and actions together.
- The feature vector of a state-action pair $(s, a)$ will be denoted as $\mathbf{x}(s, a)$ and we can think of it as a function mapping state-action pairs to some vector space, e.g. $\mathbb{R}^{d}$, i.e. $\mathbf{x}(s, a)=\left(x_{1}(s, a), x_{2}(s, a), \ldots, x_{d}(s, a)\right)^{T}$.


## Approximation of Q-Function

- Linear function approximation: Scalar product of a weight vector with the feature vector, which represents the state-action pair:

$$
\hat{Q}^{\pi}(s, a ; \mathbf{w})=\mathbf{w}^{T} \mathbf{x}(s, a)
$$

- Neural network function approximation:

$$
\hat{Q}^{\pi}(s, a ; \mathbf{w})=g(\mathbf{x}(s, a) ; \mathbf{w})
$$

where $g$ is a function represented as a neural network.

## Weight Updates

- MC:

$$
\mathbf{w}:=\mathbf{w}+\alpha \cdot\left(g_{t}-\hat{Q}\left(s_{t}, a_{t} ; \mathbf{w}\right)\right) \cdot \nabla \hat{Q}\left(s_{t}, a_{t} ; \mathbf{w}\right)
$$

- SARSA:

$$
\mathbf{w}:=\mathbf{w}+\alpha \cdot\left(r+\gamma \hat{Q}\left(s_{t+1}, a_{t+1} ; \mathbf{w}\right)-\hat{Q}\left(s_{t}, a_{t} ; \mathbf{w}\right)\right) \cdot \nabla Q\left(s_{t}, a_{t} ; \mathbf{w}\right)
$$

- Q-Learning:

$$
\mathbf{w}:=\mathbf{w}+\alpha \cdot\left(r+\gamma \max _{a \in A} \hat{Q}\left(s_{t+1}, a ; \mathbf{w}\right)-\hat{Q}\left(s_{t}, a_{t} ; \mathbf{w}\right)\right) \cdot \nabla Q\left(s_{t}, a_{t} ; \mathbf{w}\right)
$$

## Convergence

- Next time...

