


Approximate Dynamic Programming

Sanjiban Choudhury




Cornell Bowers CIS
Computer Science

When the MDP is known:

Two Fundamental Ways to
Solve for Optimal Policy

Value Iteration



$$V^*(s) = \min_a [c(s, \pi(s)) + \gamma \mathbb{E}_{s' \sim \mathcal{T}(s,a)} V^*(s')]]$$

Policy Iteration

Which one converges faster: value/policy?

0	-	10	10	10	10	10	10	10	10	10	10
1	-	10	10	10	10	10	10	10	10	10	10
2	-	10	10	10	10	10	10	10	10	10	10
3	-	10	10	10	10	10	10	10	10	10	10
4	-	10	10	10	10	10	10	10	10	10	10
5	-	10	10	10	10	10	10	10	10	10	10
6	-	10	10	10	10	10	10	10	10	10	10
7	-	10	10	10	10	10	10	10	10	10	10
8	-	10	10	10	10	10	10	10	10	10	10
9	-	10	10	10	10	10	10	10	10	10	10
		0	1	2	3	4	5	6	7	8	9

Values

0	-	x	x	x	x	x	x	x	x	x	x
1	-	x	x	x	x	x	x	x	x	x	x
2	-	x	x	x	x	x	x	x	x	x	x
3	-	x	x	x	x	x	x	x	x	x	x
4	-	x	x	x	x	x	x	x	x	x	x
5	-	x	x	x	x	x	x	x	x	x	x
6	-	x	x	x	x	x	x	x	x	x	x
7	-	x	x	x	x	x	x	x	x	x	x
8	-	x	x	x	x	x	x	x	x	x	x
9	-	x	x	x	x	x	x	x	x	x	x
		0	1	2	3	4	5	6	7	8	9

Policy



Policy converges **faster**
than the value

Can we iterate over **policies?**

Policy Iteration

Init with some policy π

Repeat forever

Evaluate policy

$$V^\pi(s) = c(s, \pi(s)) + \gamma \mathbb{E}_{s' \sim \mathcal{T}(s, a)} V^\pi(s')$$

Improve policy

$$\pi^+(s) = \arg \min_a [c(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{T}(s, a)} V^\pi(s')]$$



Init with some policy π

Iter: 0

0	-	→	→	→	→	→	→	→	→	→	↑
1	-	→	→	→	→	→	→	→	→	→	↑
2	-	→	→	→	→	→	→	→	→	→	↑
3	-	→	→	→	→	→	→	→	→	→	↑
4	-	→	→	→	→	→	→	→	→	→	↑
5	-	→	→	→	→	→	→	→	→	→	↑
6	-	→	→	→	→	→	→	→	→	→	↑
7	-	→	→	→	→	→	→	→	→	→	↑
8	-	→	→	→	→	→	→	→	→	→	↑
9	-	→	→	→	→	→	→	→	→	→	↑
		0	1	2	3	4	5	6	7	8	9

Iteration 1: Compute the value of the policy

Iter: 1



$$V^\pi(s) = c(s, \pi(s)) + \gamma \mathbb{E}_{s' \sim \mathcal{T}(s, a)} V^\pi(s')$$

$$\pi^+(s) = \arg \min_a c(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{T}(s, a)} V^\pi(s')$$

Policy Iteration

Iter: 0

0	-	0	0	0	0	0	0	0	0	0	0
1	-	0	0	0	0	0	0	0	0	0	0
2	-	0	0	0	0	0	0	0	0	0	0
3	-	0	0	0	0	0	0	0	0	0	0
4	-	0	0	0	0	0	0	0	0	0	0
5	-	0	0	0	0	0	0	0	0	0	0
6	-	0	0	0	0	0	0	0	0	0	0
7	-	0	0	0	0	0	0	0	0	0	0
8	-	0	0	0	0	0	0	0	0	0	0
9	-	0	0	0	0	0	0	0	0	0	0
		0	1	2	3	4	5	6	7	8	9

0	-	→	→	→	→	→	→	→	→	→	↑
1	-	→	→	→	→	→	→	→	→	→	↑
2	-	→	→	→	→	→	→	→	→	→	↑
3	-	→	→	→	→	→	→	→	→	→	↑
4	-	→	→	→	→	→	→	→	→	→	↑
5	-	→	→	→	→	→	→	→	→	→	↑
6	-	→	→	→	→	→	→	→	→	→	↑
7	-	→	→	→	→	→	→	→	→	→	↑
8	-	→	→	→	→	→	→	→	→	→	↑
9	-	→	→	→	→	→	→	→	→	→	↑
		0	1	2	3	4	5	6	7	8	9

$$V^\pi(s) = c(s, \pi(s)) + \gamma \mathbb{E}_{s' \sim \mathcal{T}(s, a)} V^\pi(s')$$


$$\pi^+(s) = \arg \min_a c(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{T}(s, a)} V^\pi(s')$$

When the MDP is ~~known~~ unknown:


Restricted access to the transition function

$$V^\pi(s) = c(s, \pi(s)) + \gamma \mathbb{E}_{s' \sim \mathcal{T}(s, a)} [V^\pi(s')]$$

Estimate the value of policy from sample rollouts



Roll outs



0	74	75	76	77	77	77	77	2	1	0
1	74	75	76	77	77	77	77	3	2	1
2	74	75	76	77	77	77	77	3.9	3	2
3	55	56	56	57	50	40	26	4.9	3.9	3
4	74	75	76	77	77	77	77	5.9	4.9	3.9
5	74	75	76	77	77	77	77	6.8	5.9	4.9
6	74	75	76	77	77	77	77	7.7	6.8	5.9
7	15	14	13	12	11	10	9.6	8.6	7.7	6.8
8	16	15	14	13	12	11	10	9.6	8.6	7.7
9	17	16	15	14	13	12	11	10	9.6	8.6
	0	1	2	3	4	5	6	7	8	9

Value $V^\pi(s)$



Monte-Carlo


$$V(s) \leftarrow V(s) + \alpha(G_t - V(s))$$

Zero Bias

High Variance

Always convergence

(Just have to wait till heat death of the universe)



Temporal Difference

$$V(s) \leftarrow V(s) + \alpha(c + \gamma V(s') - V(s))$$

Can have bias


Low Variance

May *not* converge if
using function approximation

Tabular setting is cute.


But how do we estimate $V(s)$
in the continuous setting





Unknown function

Researchers



Neural Network


Activity!



A *tiny* MDP

Reward for being at any state is 0.0

Discount factor $\gamma = 0.9$




What happens when you run value iteration?

(Initialize with random values, say $V(s_1) = 1$ and 2)

A *tiny* MDP

Reward for being at any state is 0.0

Discount factor $\gamma = 0.9$



Let's say we want to use a *linear value function approximator*


$$V(s) = wf(s) = w * \begin{cases} 1 & \text{if } s = s_1 \\ 2 & \text{if } s = s_2 \end{cases}$$

What happens if you run value iteration? (Initialize with $w=1$)

Think-Pair-Share

Think (30 sec): Initialize value iteration with $w=1$. What happens?
What's the explanation?

Pair: Find a partner




Share (45 sec): Partners exchange ideas

$$V(s) = wf(s)$$

Init with $w = 1$

CURSE OF APPROXIMATION!



Approximation
introduces an error that
gets amplified by both
value / policy iteration


Key separation between SL and RL

From dynamic
programming to
Fitted
dynamic
programming



Approximate (Fitted) Value Iteration

Q-iteration



$$Q(s, a) \leftarrow 0$$

while *not converged* **do**


for $s \in S, a \in A$

$$Q^{new}(s, a) = c(s, a) + \gamma \mathbb{E}_{s'} \min_{a'} Q(s', a')$$

$$Q \leftarrow Q^{new}$$

return Q

Fitted Q-iteration



Given $\{s_i, a_i, c_i, s'_i\}_{i=1}^N$

$$\text{Init } Q_\theta(s, a) \leftarrow 0$$

while *not converged* **do**

$$D \leftarrow \emptyset$$

for $i \in 1, \dots, n$

$$\text{input} \leftarrow \{s_i, a_i\}$$



$$\text{target} \leftarrow c_i + \gamma \min_{a'} Q_\theta(s'_i, a')$$

$$D \leftarrow D \cup \{\text{input}, \text{output}\}$$


$$Q_\theta \leftarrow \text{Train}(D)$$

return Q_θ


A simple example: Gridworld




What happens when we run value iteration with a *quadratic*?





What happens when we run value iteration with a
quadratic?



What happens when we run value iteration with a
quadratic?





Another Example: Mountain Car!





What happens when we run value iteration with a *2 Layer MLP?*

Iteration 11





What happens when we run value iteration with a *2 Layer MLP?*

Iteration 101



What happens when we run value iteration with a *2 Layer MLP?*



Iteration 201



The problem of Bootstrapping!



`max()`



The problem of Bootstrapping!



Errors in approximation are amplified

Key reason is the minimization

Minimization operation visit states where approximate values is less than the true value of that state – that is to say, states that look more attractive than they should.

Typically states where you have very few samples

$\max()$



What about policy iteration?



Policy Iteration

Policy Evaluation

0	-	0	0	0	0	0	0	0	0	0	0
1	-	0	0	0	0	0	0	0	0	0	0
2	-	0	0	0	0	0	0	0	0	0	0
3	-	0	0	0	0	0	0	0	0	0	0
4	-	0	0	0	0	0	0	0	0	0	0
5	-	0	0	0	0	0	0	0	0	0	0
6	-	0	0	0	0	0	0	0	0	0	0
7	-	0	0	0	0	0	0	0	0	0	0
8	-	0	0	0	0	0	0	0	0	0	0
9	-	0	0	0	0	0	0	0	0	0	0
		0	1	2	3	4	5	6	7	8	9

Iter: 0

Policy Improvement

0	-	→	→	→	→	→	→	→	→	→	↑
1	-	→	→	→	→	→	→	→	→	→	↑
2	-	→	→	→	→	→	→	→	→	→	↑
3	-	→	→	→	→	→	→	→	→	→	↑
4	-	→	→	→	→	→	→	→	→	→	↑
5	-	→	→	→	→	→	→	→	→	→	↑
6	-	→	→	→	→	→	→	→	→	→	↑
7	-	→	→	→	→	→	→	→	→	→	↑
8	-	→	→	→	→	→	→	→	→	→	↑
9	-	→	→	→	→	→	→	→	→	→	↑
		0	1	2	3	4	5	6	7	8	9

$$Q^\pi(s, a) = c(s, a) + \gamma \mathbb{E}_{s' \sim \mathcal{T}(s, a)} [Q^\pi(s', \pi(s'))]$$


$$\pi^+(s) = \arg \min_a Q^\pi(s, a)$$

Approximate (Fitted) Policy Iteration

Fitted policy evaluation

Policy Improvement

Given $\{s_i, a_i, c_i, s'_i\}_{i=1}^N$



Init $Q_\theta(s, a) \leftarrow 0$

while *not converged* **do**

$D \leftarrow \emptyset$

for $i \in 1, \dots, n$

input $\leftarrow \{s_i, a_i\}$

target $\leftarrow c_i + \gamma Q_\theta(s'_i, \pi(s'_i))$

$D \leftarrow D \cup \{\text{input, output}\}$

$Q_\theta \leftarrow \mathbf{Train}(D)$

return Q_θ


This remains
the same!

$$\pi^+(s) = \arg \min_a Q^\pi(s, a)$$


Surely approximate value
evaluation is more stable than
approximate value iteration?
(There is no `min()`!)




Well ... not quite



w blows up!




Well ... not quite



But we can fix this by
on-policy weighting

w blows
up!




Weight each datapoint
by how often the
policy visits it.

But what about policy improvement?

Fitted policy evaluation

Policy Improvement

Given $\{s_i, a_i, c_i, s'_i\}_{i=1}^N$



Init $Q_\theta(s, a) \leftarrow 0$

while *not converged* do

This is fine..
 $D \leftarrow \emptyset$

for $i \in 1, \dots, n$

input $\leftarrow \{s_i, a_i\}$

target $\leftarrow c_i + \gamma Q_\theta(s'_i, \pi(s'_i))$

$D \leftarrow D \cup \{\text{input, output}\}$

$Q_\theta \leftarrow \text{Train}(D)$

return Q_θ


But this has
the $\min()$ step!

$$\pi^+(s) = \arg \min_a Q^\pi(s, a)$$

The problem of distribution shift

Upper half of state
is BAD

Lower half of state
is GOOD



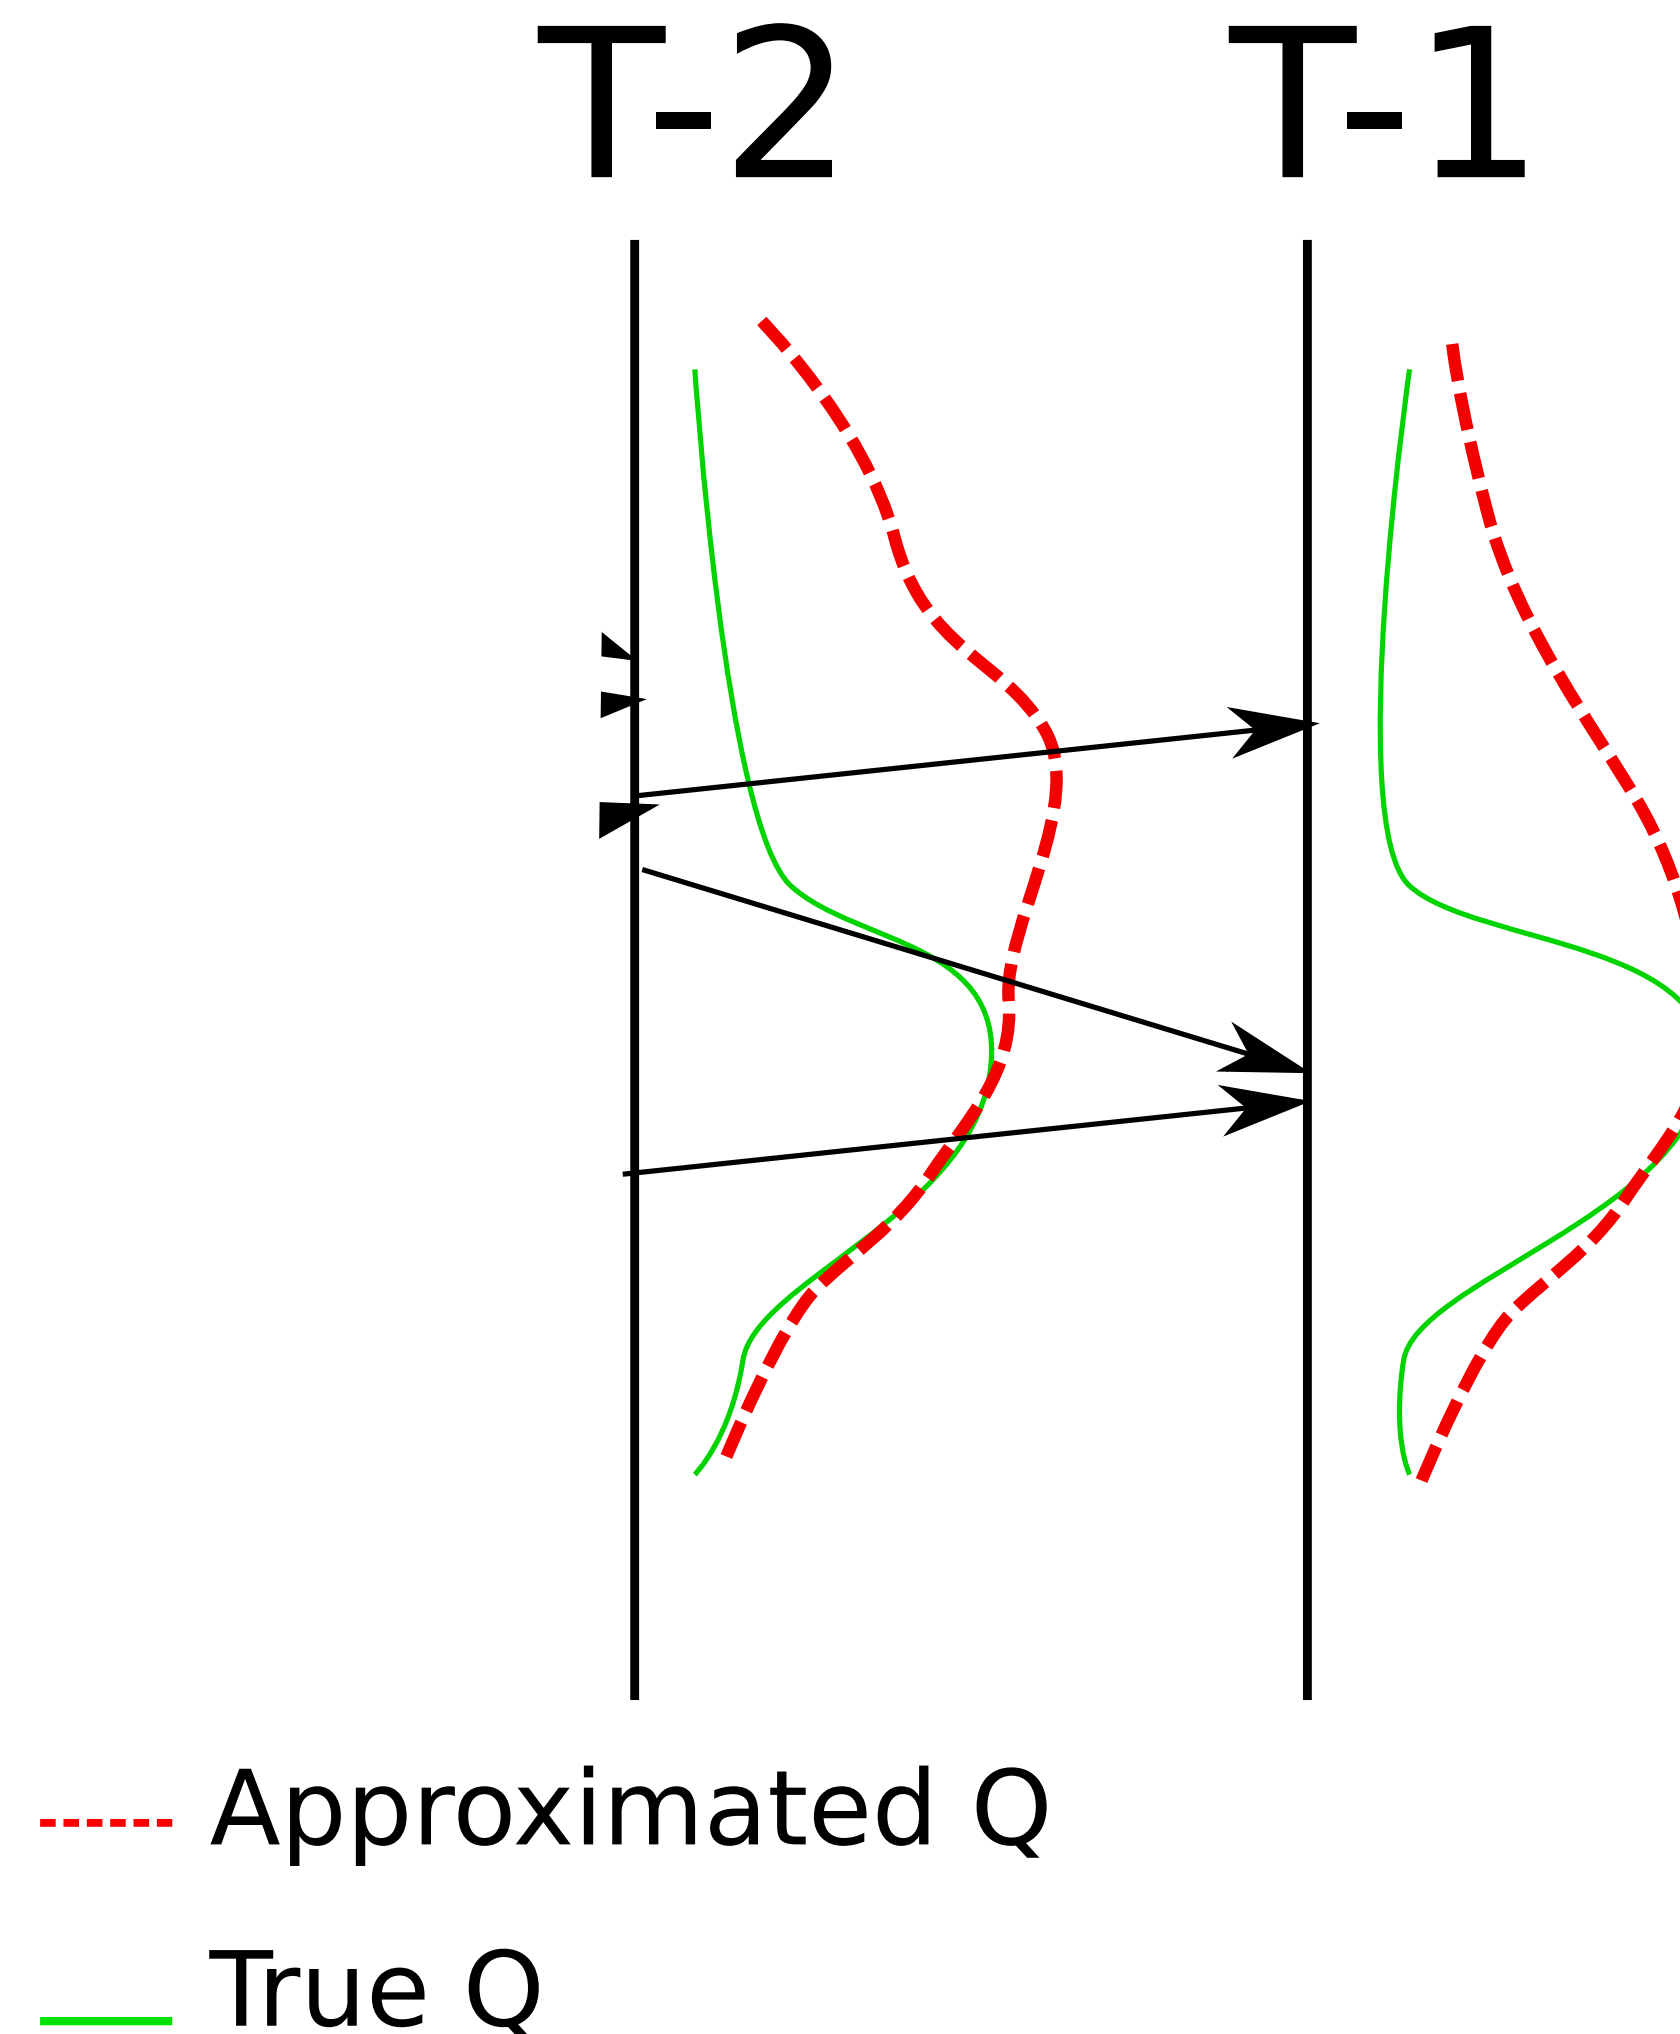
----- Approximated Q

— True Q


The problem of distribution shift

Upper half of state
is BAD

Lower half of state
is GOOD



The problem of distribution shift




Bootstrapping

Distribution Shift




Two sides of the same coin

Bootstrapping




Distribution shift




Ideas for fixing
this?



Remedies



Bootstrapping



When doing $\min()$,
don't trust value estimate

Execute policy
and **trust actual returns**

Distribution shift




Minimize the
distribution shift


Be conservative,
change policy slowly



Remedies



Bootstrapping



When doing $\min()$,
don't trust value estimate

Execute policy
and **trust actual returns**

Distribution shift




Minimize the
distribution shift

Be conservative,
change policy slowly

tl;dr

Approximate (Fitted) Value Iteration

Q-iteration




```

 $Q(s, a) \leftarrow 0$ 
while not converged do
  for  $s \in S, a \in A$ 
     $Q^{new}(s, a) = c(s, a) + \gamma \mathbb{E}_{s'} \min_{a'} Q(s', a')$ 
   $Q \leftarrow Q^{new}$ 
return  $Q$ 

```

Fitted Q-iteration



Given $\{s_i, a_i, c_i, s'_i\}_{i=1}^N$

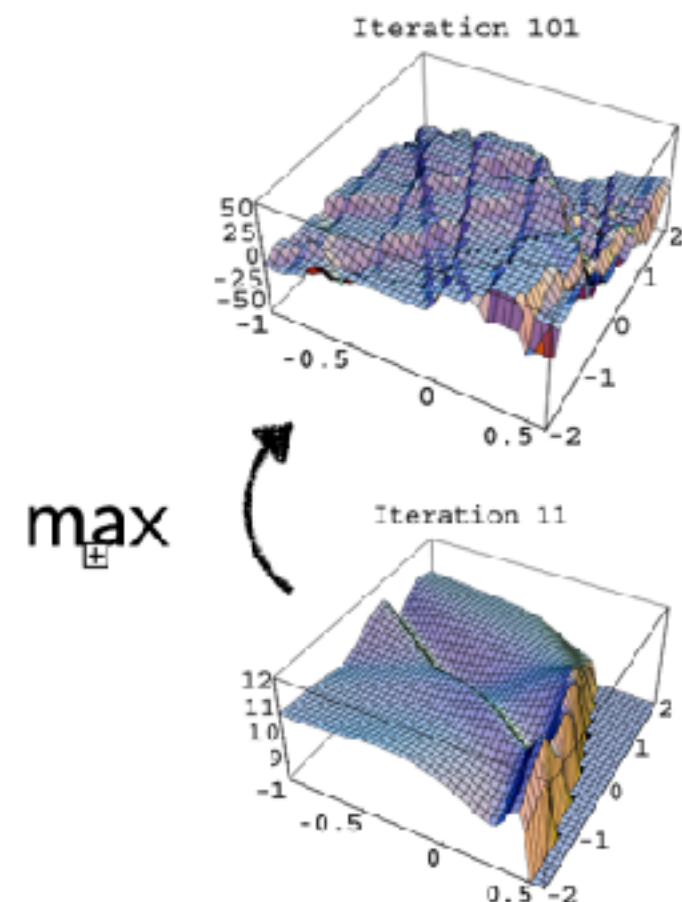
```

Init  $Q_\theta(s, a) \leftarrow 0$ 
while not converged do
   $D \leftarrow \emptyset$ 
  for  $i \in 1, \dots, n$ 
    input  $\leftarrow \{s_i, a_i\}$ 
    target  $\leftarrow c_i + \gamma \min_{a'} Q_\theta(s'_i, a')$ 
     $D \leftarrow D \cup \{\text{input}, \text{output}\}$ 
   $Q_\theta \leftarrow \text{Train}(D)$ 
return  $Q_\theta$ 


```

1

Bootstrapping



Distribution shift



Remedies



Bootstrapping



When doing $\min()$,
don't trust value estimate

Execute policy
and **trust actual returns**

Distribution shift



Minimize the
distribution shift

Be conservative,
change policy slowly

