

Deep Q Network (DQN)

Announcements

We will release HW2 tonight (Q-learning, TD,
and simulation lemma)

We will release the first reading quiz today

Recap: Q-learning

Tabular Q Learning: maintain a table \hat{Q} of size $S \times A$

$$\hat{Q}(s, a) \leftarrow \hat{Q}(s, a) + \eta \left(r + \gamma \max_{a'} \hat{Q}(s', a') - \hat{Q}(s, a) \right)$$

Data collection via ϵ -greedy:

W/ prob ϵ , select action uniform randomly

W/ prob $1 - \epsilon$, select greedy action $\arg \max_a \hat{Q}(s, a)$

Today

Consider large-scale MDPs,

how to estimate $Q^*(s, a)$ using function approximation (e.g., neural network)

**Deep Q-network (DQN) is the earliest example
of showing Deep Learning + RL is powerful**

doi:10.1038/nature14236

Human-level control through deep reinforcement learning


Volodymyr Mnih^{1*}, Koray Kavukcuoglu^{1*}, David Silver^{1*}, Andrei A. Rusu¹, Joel Veness¹, Marc G. Bellemare¹, Alex Graves¹, Martin Riedmiller¹, Andreas K. Fidjeland¹, Georg Ostrovski¹, Stig Petersen¹, Charles Beattie¹, Amir Sadik¹, Ioannis Antonoglou¹, Helen King¹, Dharshan Kumaran¹, Daan Wierstra¹, Shane Legg¹ & Demis Hassabis¹

Outline:

1. Q Learning w/ function approximation
2. Replay buffer, batch optimization and target network

Q-Learning w/ function approximation

We will model Q^* using a function approximator



$$Q_{\theta}(s, a) : S \times A \rightarrow \mathbb{R}$$

Assumption: differentiable $\nabla_{\theta} Q_{\theta}(s, a)$

(The DQN paper uses ConvNet as s is an image frame of the game)

$$x = [s^{\top}, a, 1]^{\top}$$
$$y = \alpha^{\top} \text{ReLU} \left(W^{[2]} \text{ReLU} \left(W^{[1]} x \right) \right) + b$$

Attempt 1: Q Learning w/ function approximation

Initialize θ^0 . Set initial state $s \in \mathcal{S}$

For $t = 0$ to T

Take action a based on ϵ -greedy of Q_{θ_t} , get reward r and next state $s' \sim P(\cdot | s, a)$

Form Q-target $r + \gamma \max_{a'} Q_{\theta_t}(s', a')$

Update to θ^{t+1} :

Set $s \leftarrow s'$

Q Learning w/ function approximation

Update parameters using SGD on the Bellman error loss

$$\ell_{be}(\theta) := (Q_{\theta}(s, a) - y)^2, \text{ where } y = r(s, a) + \gamma \mathbb{E}_{s' \sim P(\cdot | s, a)} \max_{a'} Q_{\theta}(s', a')$$

Issues of this simple approach

1. Inefficient — it throws away all historical data
(your network could forget old experiences, i.e., catastrophic forgetting)

2. instability — Training is quite unstable (we saw it from the past Cartpole Demo)

Outline:

1. Q Learning w/ function approximation
2. Replay buffer, batch optimization and target network

Q-Learning w/ function approximation

First improvement: Replay buffer

$$\mathcal{D}_{rb} = \begin{bmatrix} \dots \\ (s, a, r, s') \\ \dots \\ \dots \end{bmatrix}$$

A dataset that contains all historical
State-action-reward-next state tuples

Q-Learning w/ function approximation

With replay buffer, we can use **mini-batch SGD** to update Bellman error loss

Given Q_{θ^t} and replay buffer \mathcal{D}_{rb} , randomly sample a mini-batch \mathcal{B} from \mathcal{D}_{rb}

$$\theta^{t+1} = \theta^t - \eta \frac{1}{|\mathcal{B}|} \sum_{(s,a,r,s') \in \mathcal{B}} \left(Q_{\theta^t}(s, a) - r - \max_{a'} Q_{\theta^t}(s', a') \right) \nabla_{\theta} Q_{\theta^t}(s, a)$$

Q-Learning w/ function approximation

Second improvement: Target network (making Q learning more stable)

Recall that Q learning can be understood as running SGD on an **evolving** loss function

$$\ell_{be}(\theta) := \left(Q_{\theta}(s, a) - y \right)^2, \text{ where } y = r(s, a) + \underbrace{\gamma \mathbb{E}_{s' \sim P(\cdot | s, a)} \max_{a'} Q_{\theta}(s', a')}_{\text{regression target}}$$

Source of instability: target changes immediately whenever we update θ

Q-Learning w/ function approximation

Second improvement: Target network (making Q learning more stable)

Introducing target network $Q_{\tilde{\theta}}$ to **slow down** the evolution of the BE loss

(e.g., set $\tilde{\theta}$ as a copy of an older version of θ)

$$\ell_{be}(\theta) := (Q_{\theta}(s, a) - y)^2, \text{ where } y = r(s, a) + \gamma \mathbb{E}_{s' \sim P(\cdot | s, a)} \max_{a'} Q_{\tilde{\theta}}(s', a')$$

Use a target network
that slowly catches up θ

Attempt 2: Deep Q network (DQN)

Initialize θ and replay buffer \mathcal{D}_{rb} . Set $\tilde{\theta} = \theta$, Set initial state $s \in \mathcal{S}$

While true:

Take action a based on ϵ -greedy of Q_θ , get reward r and next state $s' \sim P(\cdot | s, a)$

Add (s, a, r, s') to \mathcal{D}_{rb}

Sample mini-batch \mathcal{B} from \mathcal{D}_{rb}

Update parameters:

$$\theta \leftarrow \theta - \eta \frac{1}{|\mathcal{B}|} \sum_{(s,a,r,s') \in \mathcal{B}} \left(Q_\theta(s, a) - r - \max_{a'} Q_{\tilde{\theta}}(s', a') \right) \nabla_\theta Q_\theta(s, a)$$

Every C step, set $\tilde{\theta} = \theta$

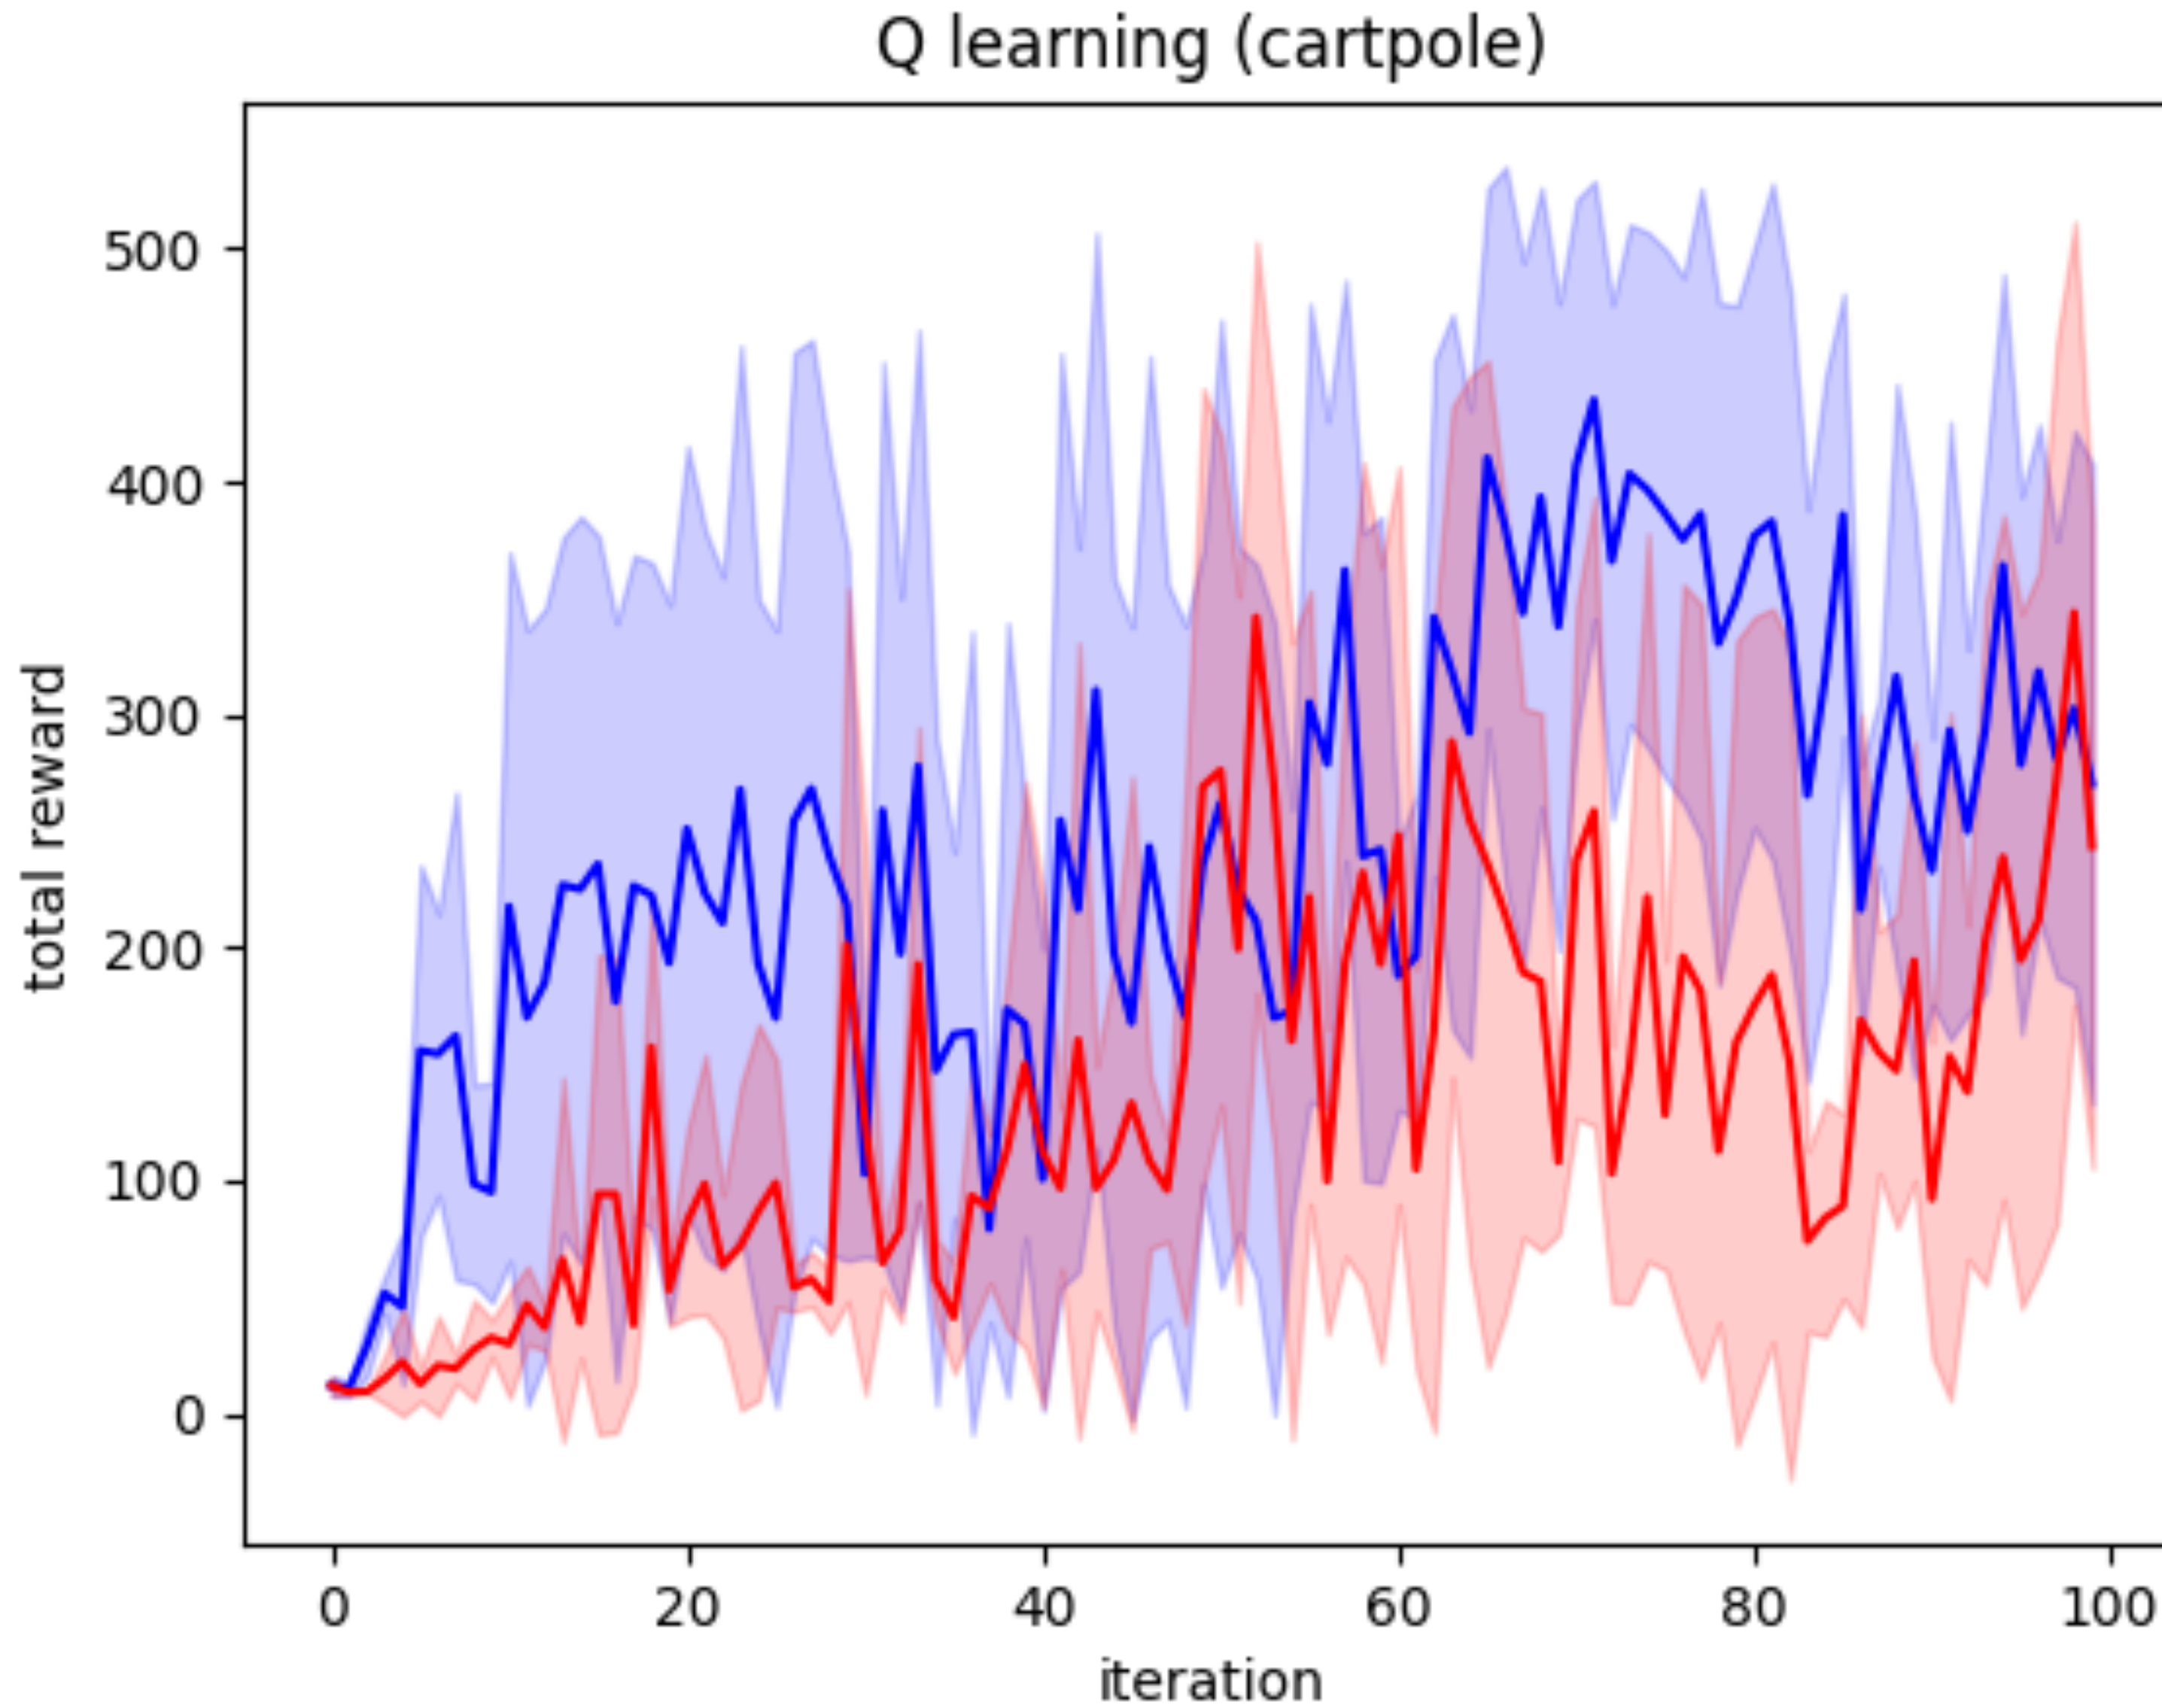
When C is large...

DQN is performing standard regression between two target updates..

$$\min_{\theta} \sum_{s,a,r,s' \in \mathcal{D}_{rb}} \left(Q_{\theta}(s, a) - \left(r + \max_{a'} Q_{\tilde{\theta}}(s', a') \right) \right)^2$$

Q: What is the Bayes optimal of this regression problem?

DQN vs Naive Q-learning



BLUE: DQN

RED: Q-learning

Summary

1. Using function approximation to handle large state space
2. Making Q-learning closer to the supervised learning (i.e., regression) framework:
 - Replay buffer + mini-batch SGD
 - Target network to simulate a standard regression setting