



Practical Machine Learning

Workshop 5. Reinforcement Learning (RL)

Dr. Suyong Eum

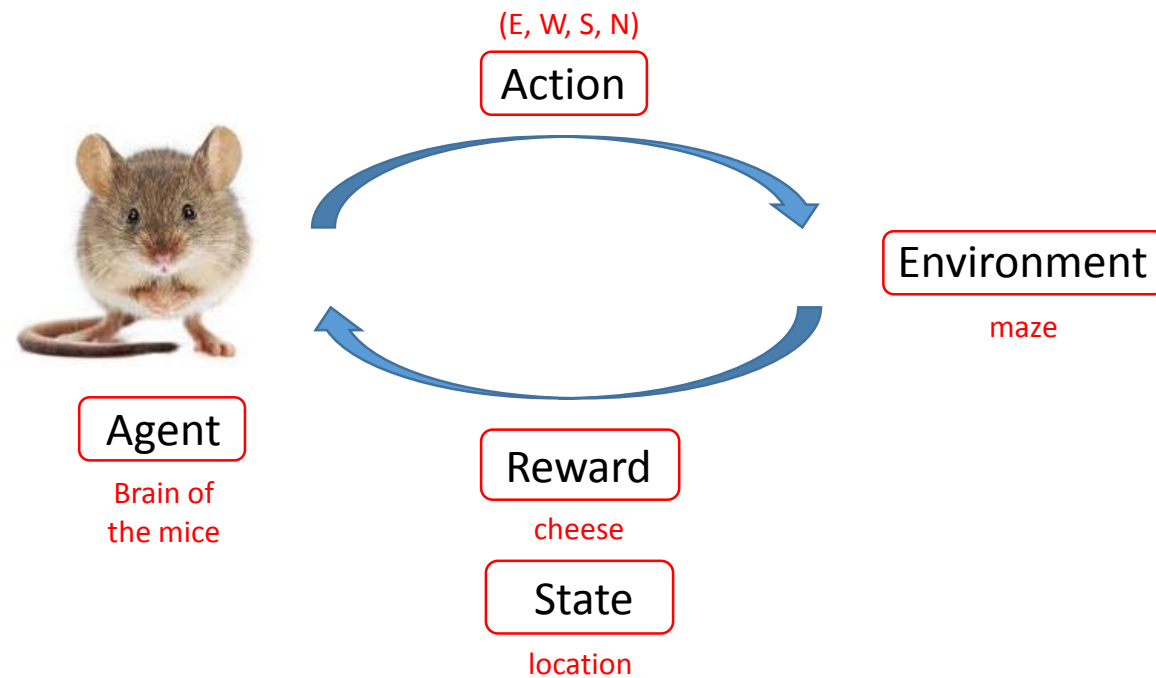


You are going to learn

- ❑ What Reinforcement Learning (RL) is
- ❑ Deep Q Network (DQN)
- ❑ Policy Gradient (PG)
 - Actor Critic (AC)
 - Asynchronous Advantage Actor Critic (A3C)
- ❑ Hand-on experience
 - CartPole game

Big picture: Reinforcement Learning (RL)

- ❑ Learning how to **take actions** in an **environment** so as to maximize **future cumulative reward**.



Reinforcement Learning (RL) has been around for a long time

- ❑ Reinforcement Learning (RL) has been used for many applications where an agent interacts with an environment while trying **to learn optimal sequence of decisions** – optimal control problems:
 - Manufacturing, e.g., robot arms to assemble cars.
 - Financial strategy, e.g., buy or sell to maximize the value of the portfolio
 - Inventory management or resource management

- Crites, R.H. and Barto, A.G. (1998). Elevator Group Control Using Multiple Reinforcement Learning Agents. Machine Learning, 33:235-262.
- Kaelbling, L. P., Littman, M. L., and Moore, A. W. (1996). Reinforcement learning: A Survey. Journal of Artificial Intelligence Research, 4:237-285.
- Singh, S., Kearns, M. (2002). Near-Optimal Reinforcement Learning in Polynomial Time. Machine Learning journal, Volume 49, Issue 2, pages 209-232, 2002.
- Schultz, W. (2002). Getting formal with dopamine and reward. Neuron, 36:241-263.
- Sutton, R. S. (1984). Temporal credit assignment in reinforcement learning. PhD thesis, University of Massachusetts, Amherst
- Sutton, R. S. (1996). Generalization in reinforcement learning: Successful examples using sparse coding. In D. S. Touretzky, M. C. Mozer and M. E. Hasselmo (eds.) Advances in Neural Information Processing, pp. 1038-1044, MIT Press, Cambridge, CA.
- Sutton, R. S. and Barto, A. G. (1998). Reinforcement Learning: An Introduction. Bradford Books, MIT Press, Cambridge, MA, 2002 edition.



Shed new light on RL with neural networks

Playing Atari with Deep Reinforcement Learning

Volodymyr Mnih Koray Kavukcuoglu David Silver Alex Graves Ioannis Antonoglou

Daan Wierstra Martin Riedmiller

DeepMind Technologies

{vlad,koray,david,alex.graves,ioannis,daan,martin.riedmiller}@deepmind.com

Abstract

We present the first deep learning model to successfully learn control policies directly from high-dimensional sensory input using reinforcement learning. The model is a convolutional neural network, trained with a variant of Q-learning, whose input is raw pixels and whose output is a value function estimating future rewards. We apply our method to seven Atari 2600 games from the Arcade Learning Environment, with no adjustment of the architecture or learning algorithm. We find that it outperforms all previous approaches on six of the games and surpasses a human expert on three of them.

2013

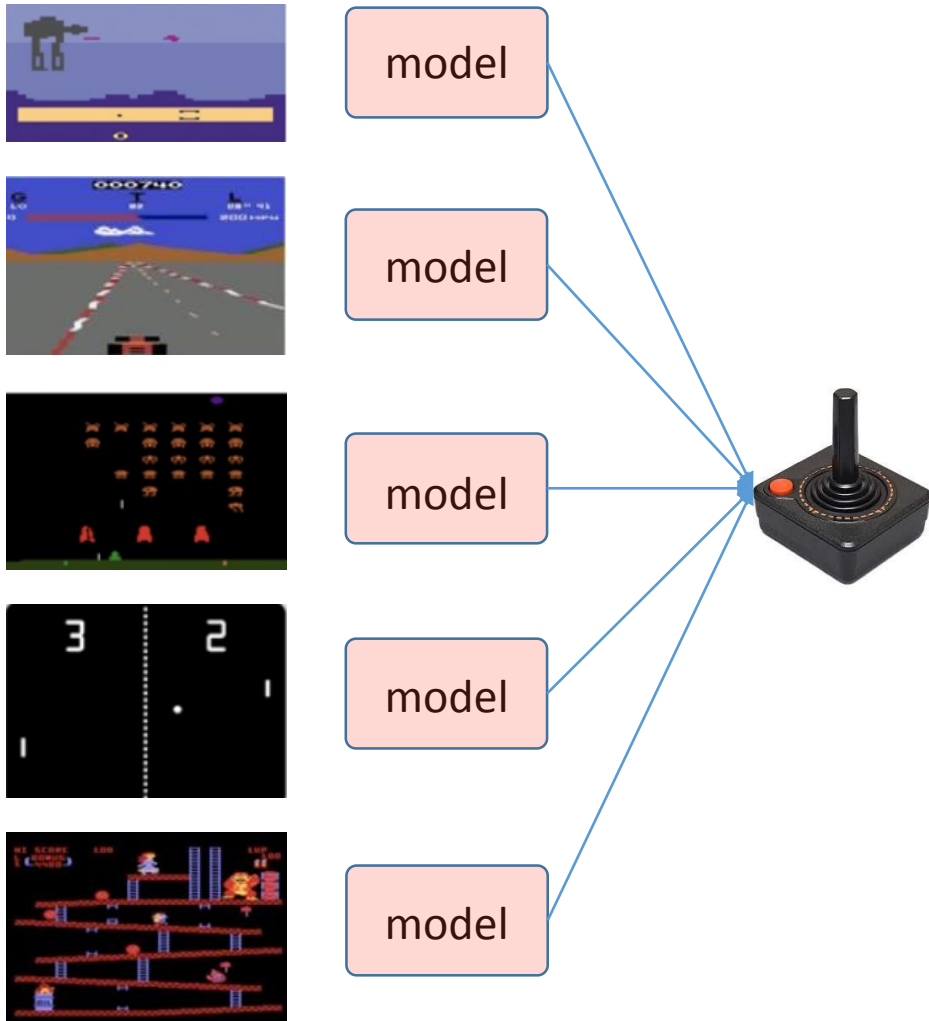


2015




2016

Shed new light on RL with neural networks - cont .




How to model the environment?

Shed new light on RL with neural networks - cont .



How to model the environment?

How good is the reinforcement learning algorithm?




20 July 2016

- Human-level control through deep reinforcement learning, Feb. 26, 2015, Nature
- <https://deepmind.com/applied/deepmind-for-google/>

Terminology

- 1) State (S)
- 2) Reward (R)
 - Discount factor (γ)
- 3) Policy (π)
 - Action (A)
- 4) Environment
 - Transit Probability (P)

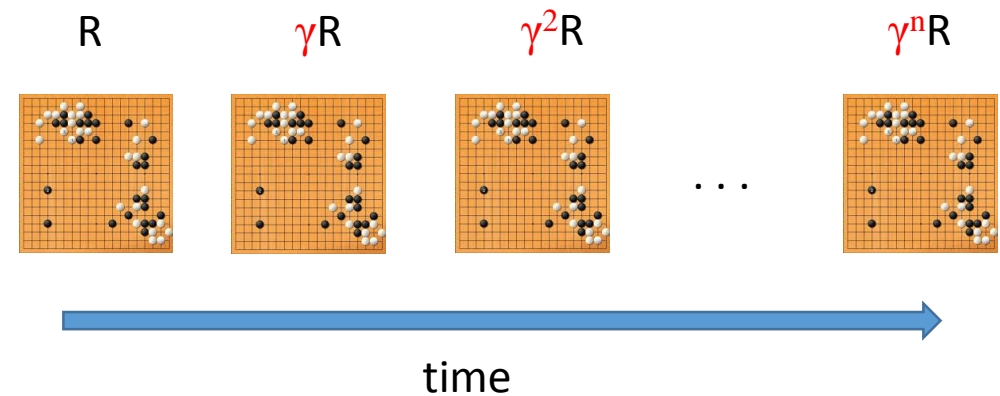
- A representation of environment that an agent recognizes.
- A agent takes an action given state
- E.g., pixel information as shown below



Reward (R) and Discount factor (γ)

- 1) State (S)
- 2) Reward (R)
 - Discount factor (γ)
- 3) Policy (π)
 - Action (A)
- 4) Environment
 - Transit Probability (P)

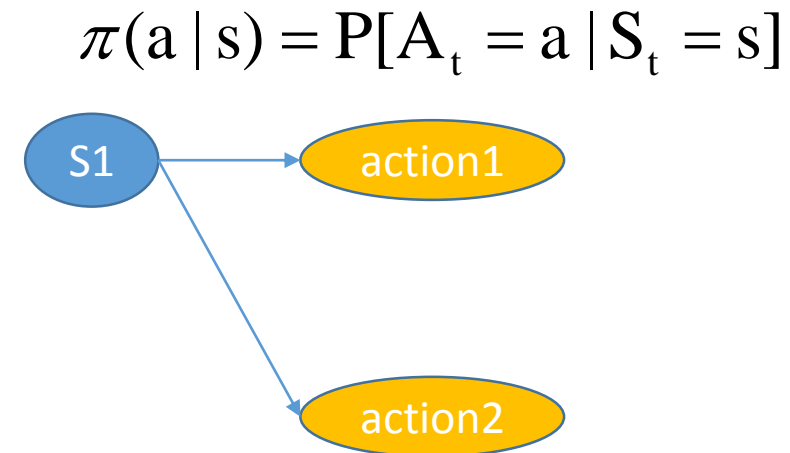
- When an agent takes an action, it considers two types of rewards
 1. Immediate reward
 2. Future accumulated reward
- Reward in a different time step may need to be treated differently: $\gamma[0,1]$



Policy (π) and Action (A)

- 1) State (**S**)
- 2) Reward (**R**)
 - Discount factor (γ)
- 3) Policy (**π**)
 - Action (**A**)
- 4) Environment
 - Transit Probability (**P**)

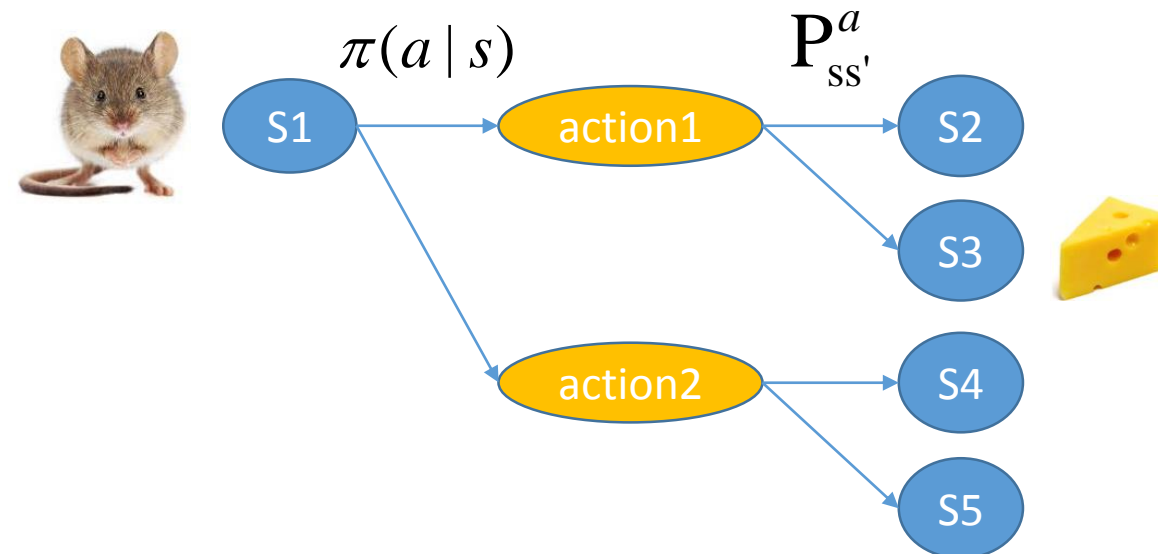
- Agent has a set of actions given state
- Policies π is a distribution over actions given states: probability of action a given state s
 - *Deterministic policy*: $a = \pi(s)$
 - *Stochastic policy*: $\pi(a|s)$



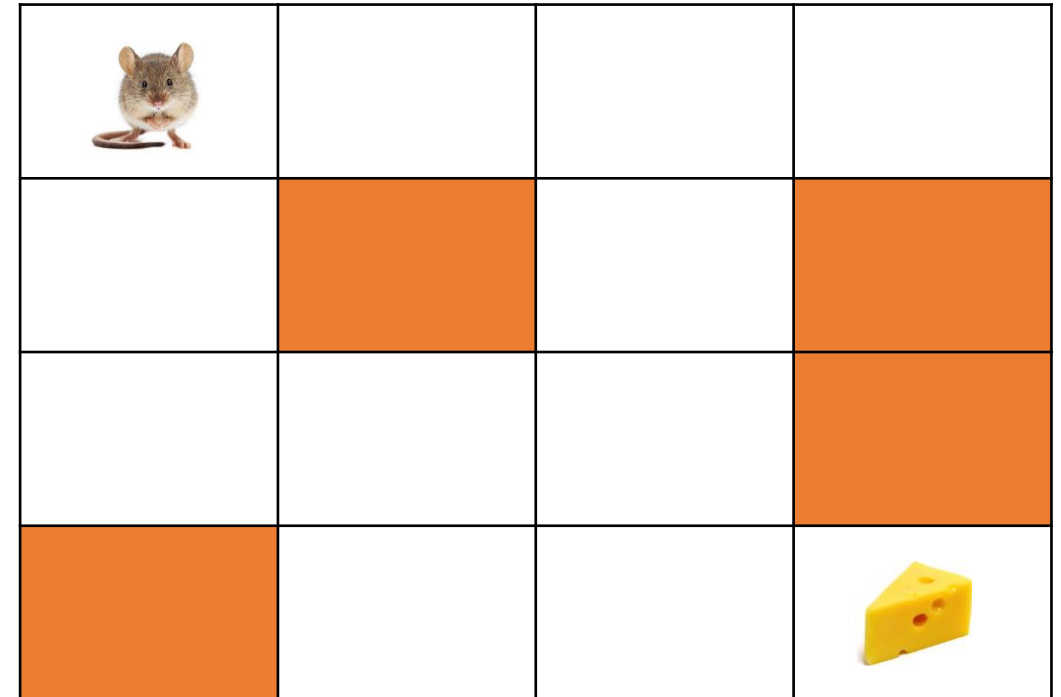

- 1) State (**S**)
- 2) Reward (**R**)
 - Discount factor (γ)
- 3) Policy (π)
 - Action (**A**)
- 4) Environment
 - Transit Probability (**P**)

- Due to uncertainty of environment, an action taken by an agent does not guarantee to a certain state.
- The uncertainty is represented as transition probability.


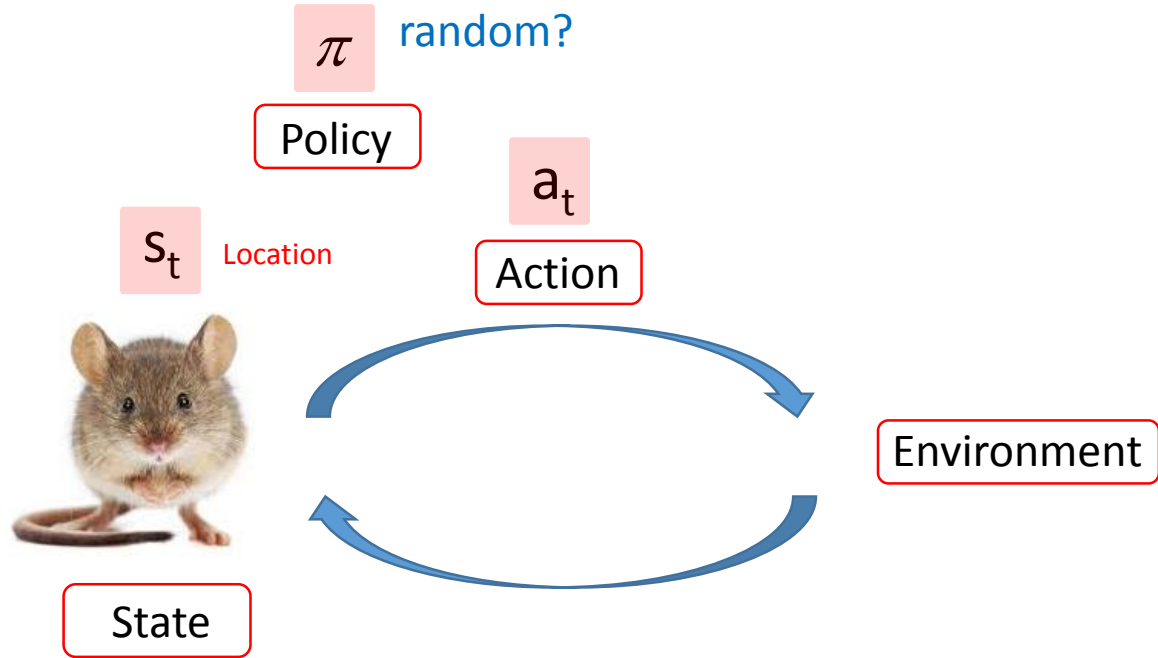
$$P_{ss'}^a = P^a[S_{t+1} = s' | S_t = s, A_t = a]$$




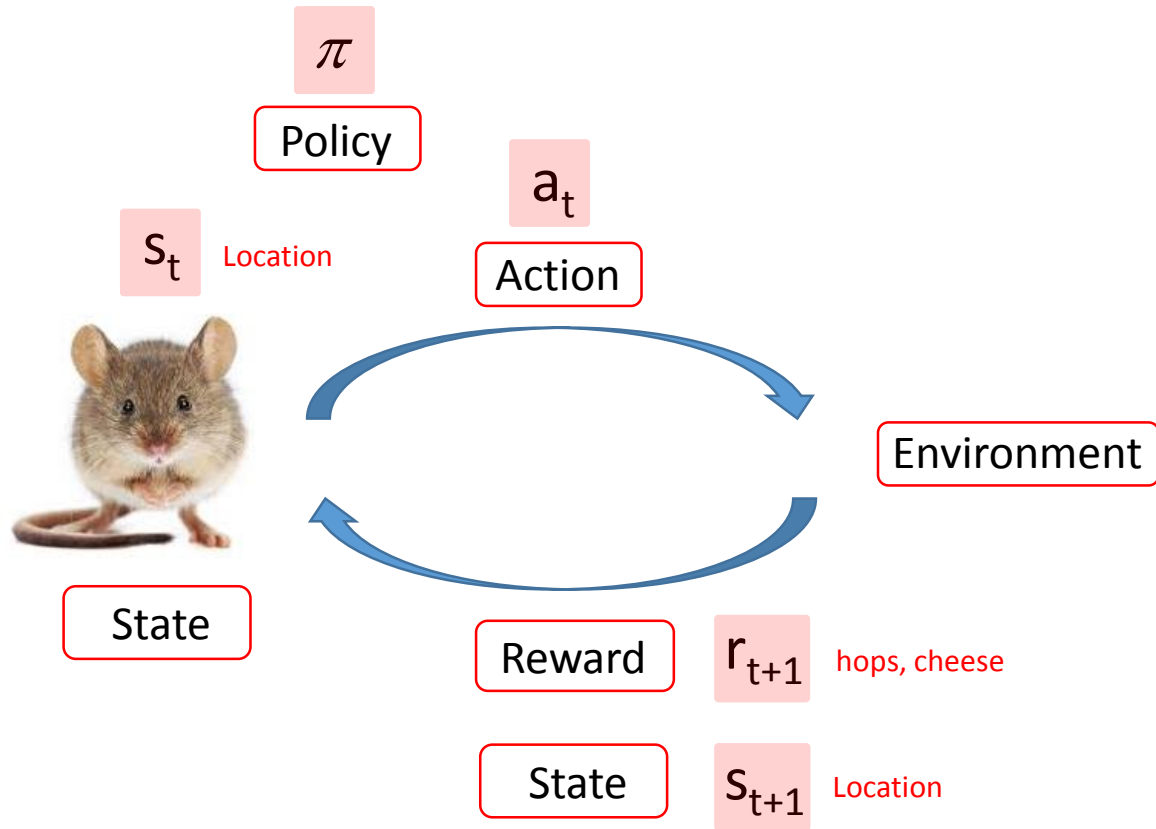
Reinforcement Learning (RL) process



Reinforcement Learning (RL) process



Reinforcement Learning (RL) process



$$r_t = r_{t+1} + \gamma r_{t+2} + \gamma^2 r_{t+3} + \dots = \sum_{k=1}^{\infty} \gamma^{k-1} r_{t+k}$$


How to find the optimum policy?: two functions as a signalpost

1) State value function: $V_{\pi}(s)$

- Expected reward when starting in state s and following policy π thereafter.
- The mouse needs to know the value of next state before making an action.

2) Action value function: $Q_{\pi}(s,a)$

- Expected reward when taking action a in state s and following policy π thereafter.
- The mouse just needs to take an action based on Q value.





How to find the optimum policy?: two functions as a signalpost

1) State value function: $V_{\pi}(s)$

- Expected reward when starting in state s and following policy π thereafter.
- The mouse needs to know the value of next state before making an action.

2) Action value function: $Q_{\pi}(s,a)$

- Expected reward when taking action a in state s and following policy π thereafter.
- The mouse just needs to take an action based on Q value.

| | | | |
|---|-----------------|--|---|
|  | $V_{\pi}(s)=10$ | | |
| $V_{\pi}(s)=7$ | | | |
| | | | |
| | | |  |

How to find the optimum policy?: two functions as a signalpost

1) State value function: $V_{\pi}(s)$


- Expected reward when starting in state s and following policy π thereafter.
- The mouse needs to know the value of next state before making an action.

2) Action value function: $Q_{\pi}(s,a)$



- Expected reward when taking action a in state s and following policy π thereafter.
- The mouse just needs to take an action based on Q value.

$Q_{\pi}(s,a1) ?$

$Q_{\pi}(s,a2) ?$



How to find the $Q(s,a)$ function?: Q learning



$$Q_{\pi}(s, a) = r_s^a + \gamma \max_{a'} Q_{\pi}(s', a')$$


Immediate
reward

Discounted reward
of successor state

Q learning: example


Q learning: $Q(s,a)$ based on Q table

- ❑ One box represents a state and actions which can be taken by an agent (N, E, S, W)
- ❑ Initial Q values at individual states are set to zero




Q learning: $Q(s,a)$ based on Q table

- ❑ Agent takes an action at each state based on value of $Q(s,a)$.
- ❑ Assuming that an agent is at the next to the goal and takes an action to E (East).



Q learning: $Q(s,a)$ based on Q table


- $Q(s, E)$ is set to one: $1 : (\text{immediate reward}) + 0 : Q(s', a')$



$$Q(s,a) = 1 : (\text{immediate reward}) + 0 : Q(s',a')$$


Q learning: $Q(s,a)$ based on Q table

- ❑ Again assume that an agent ends up to the state below and takes an action to E
- ❑ Then, $Q(s, E)$ is set to one: 0 (immediate reward) + $1 Q(s', a')$



Q learning: $Q(s,a)$ based on Q table


- ❑ Again assume that an agent ends up to the state below and takes an action to E
- ❑ Then, $Q(s, E)$ is set to one: 0 (immediate reward) + $1 Q(s', a')$



$$Q(s,a) = 0 : (\text{immediate reward}) + 1 : Q(s',a')$$

Q learning: $Q(s,a)$ based on Q table

- ❑ The route to the goal is not an optimum.
- ❑ How to select a different route occasionally?



E-greedy policy

$e = 0.1$

if $\text{rand} < e$: → 10% random decision
 action = random
else: → 90% deterministic decision
 action = $\text{argmax}(Q(s,a))$

Decaying E-greedy policy

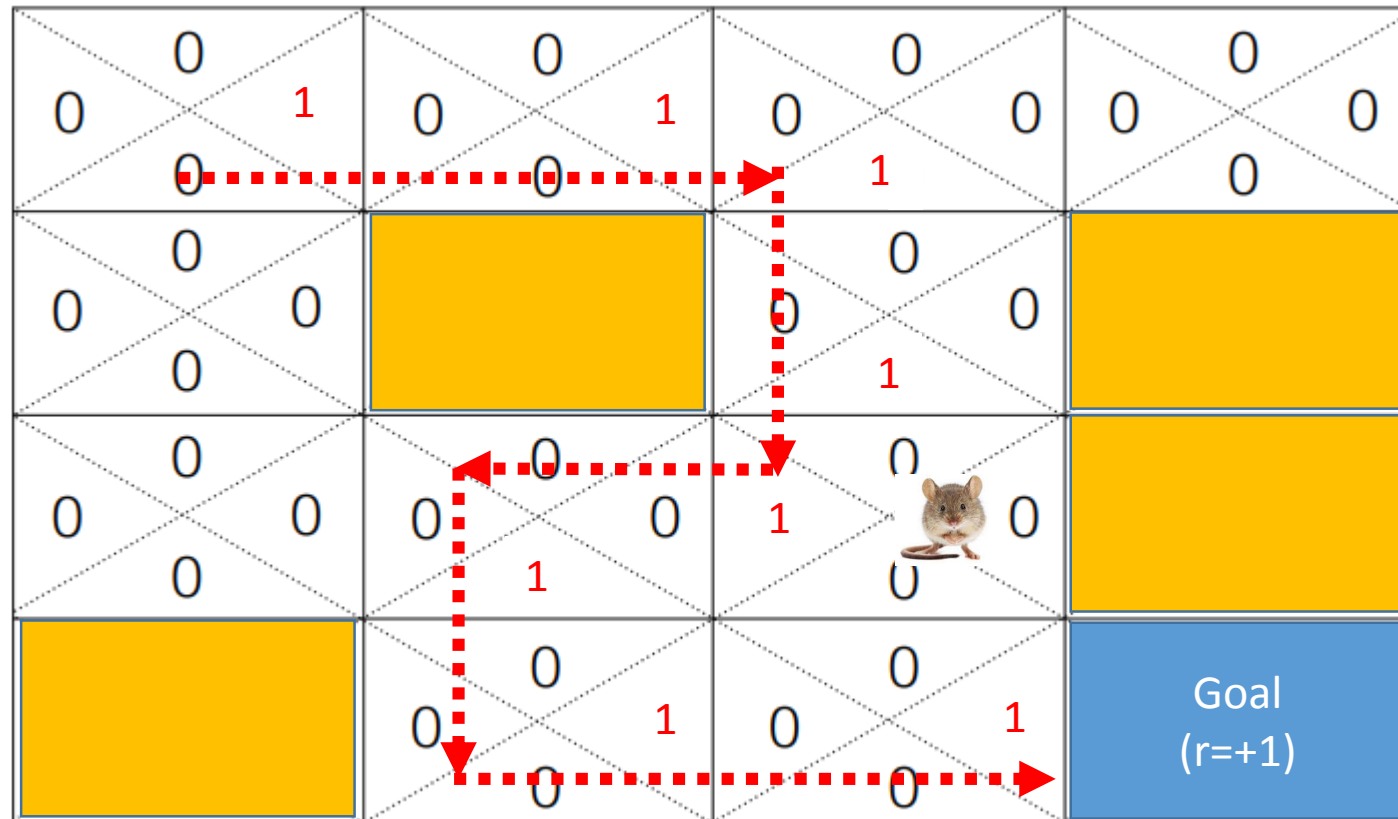
for i in range (1000)

$e = 0.1 / (i+1)$ → Random to deterministic
 decision as iteration goes on

if $\text{rand} < e$:
 action = random
else:
 action = $\text{argmax}(Q(s,a))$

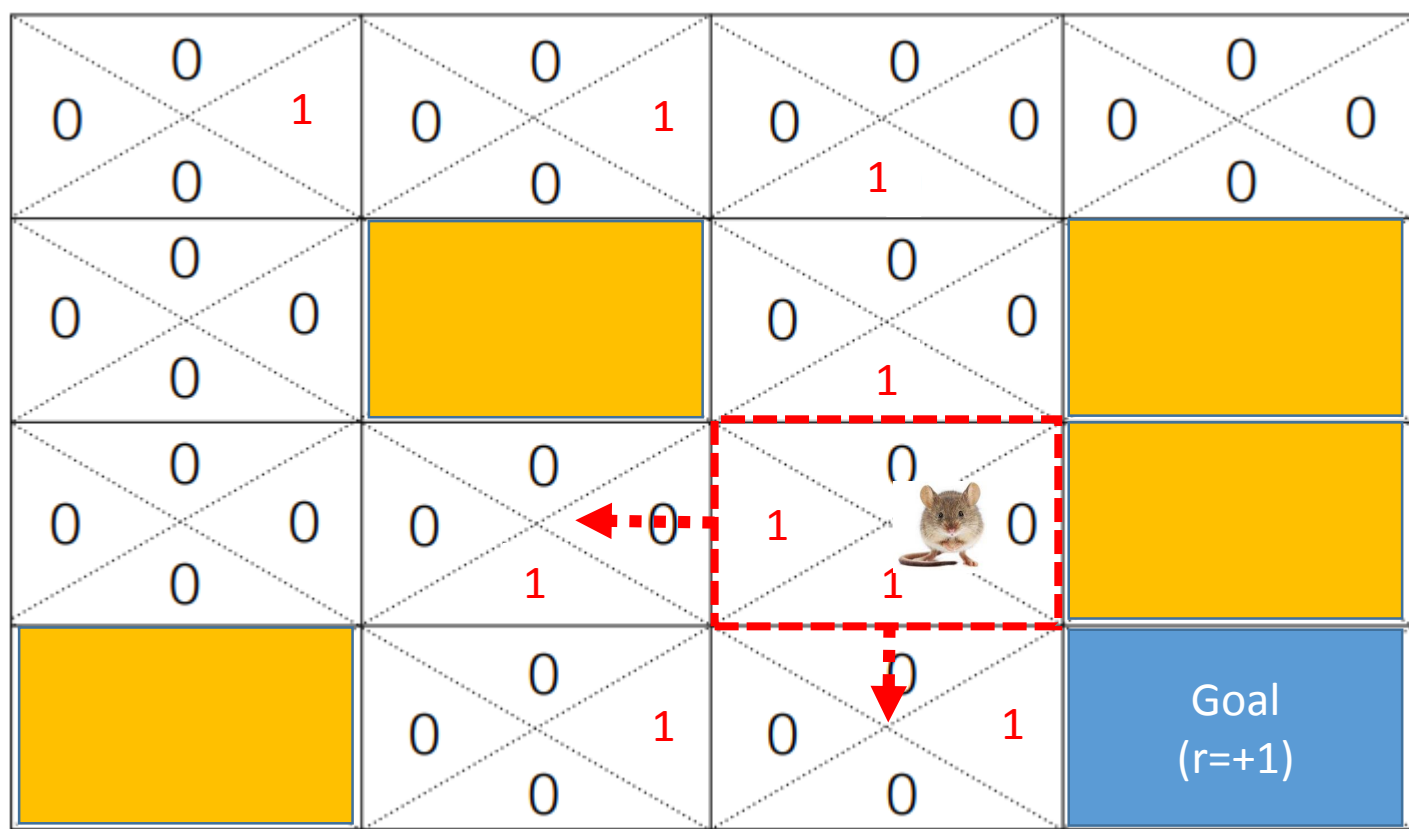
Q learning: Exploit vs Exploration

- ❑ The route to the goal is not an optimum.
- ❑ How to select a different route occasionally?



Q-learning: Discounted reward γ


- ❑ Q value does not tell which path is better



$$Q(s,a) = 0 : (\text{immediate reward}) + 1 : Q(s',a')$$

Q-learning: Discounted reward γ

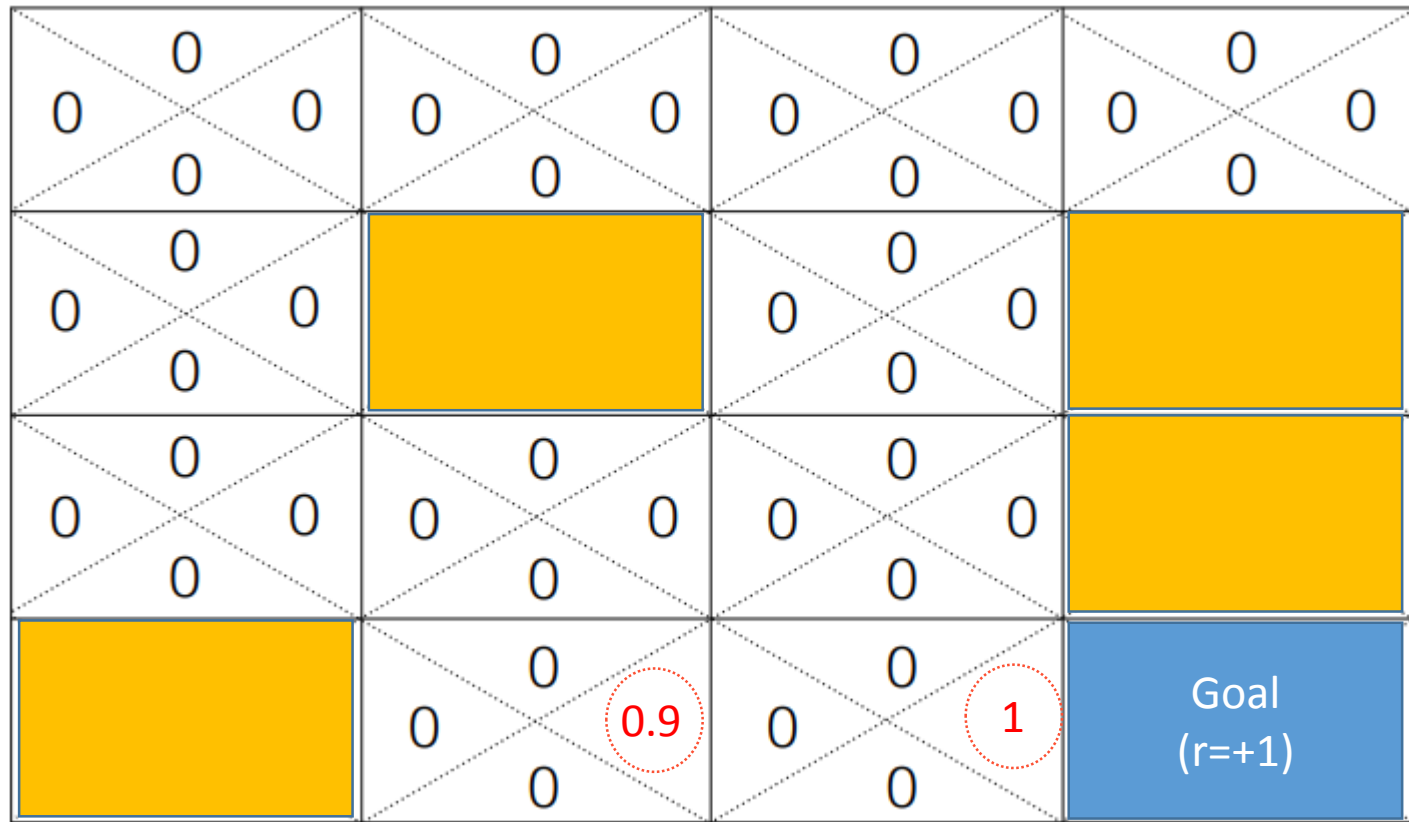
- ❑ Q value does not tell which path is better
- ❑ Let's introduce discounted reward $\gamma=0.9$ to the equation



$$Q(s,a) = r + \gamma Q(s',a') = 1 + 0.9 \times 0 = 1$$

Q-learning: Discounted reward γ


- ❑ Q value does not tell which path is better
- ❑ Let's introduce discounted reward $\gamma=0.9$ to the equation



$$Q(s,a) = r + \gamma Q(s',a') = 0 + 0.9 \times 1 = 0.9$$

Q-learning: Discounted reward γ


- ❑ Q value does not tell which path is better
- ❑ Let's introduce discounted reward $\gamma=0.9$ to the equation



$$Q(s,a) = r + \gamma Q(s',a') = 0 + 0.9 \times 0.9 = 0.81$$

Q-learning: Discounted reward γ


- ❑ Q value does not tell which path is better
- ❑ Let's introduce discounted reward $\gamma=0.9$ to the equation



$$Q(s,a) = r + \gamma Q(s',a') = 0 + 0.9 \times 0.81 = 0.729$$

Q-learning: Discounted reward γ


- ❑ Q value does not tell which path is better
- ❑ Let's introduce discounted reward $\gamma=0.9$ to the equation



$$Q(s,a) = r + \gamma Q(s',a') = 0 + 0.9 \times 1 = 0.9$$

Q-learning: Discounted reward γ

- ❑ Q value does not tell which path is better
- ❑ Let's introduce discounted reward $\gamma=0.9$ to the equation



$$Q(s,a) = r + \gamma Q(s',a') = 0 + 0.9 \times 1 = 0.9$$

For each (s, a) pair, initialize table entry $\hat{Q}(s, a)$ to zero.

Observe the current state s

Do forever:

- Select an action a and execute it
- Receive immediate reward r
- Observe the new state s'
- Update the table entry for $\hat{Q}(s, a)$ as follows:

Exploit & Exploration

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


Discounted reward

- $s \leftarrow s'$


Deep Q-Networks (DQN)

Q learning problem

- ❑ $Q(s,a)$ needs to be computed for every state(s)-action(a) pair.
 - If the problem size is small, we can handle it using Q table
- ❑ When state space is huge: computationally infeasible for entire state space
 - Backgammon: 10^{20} states
 - Computer Go: 10^{170} states
 - Automatic driving: continuous state space
- ❑ Too many states to store in memory and also too slow to learn state space



Backgammon





Go

Deep Q-Network architecture

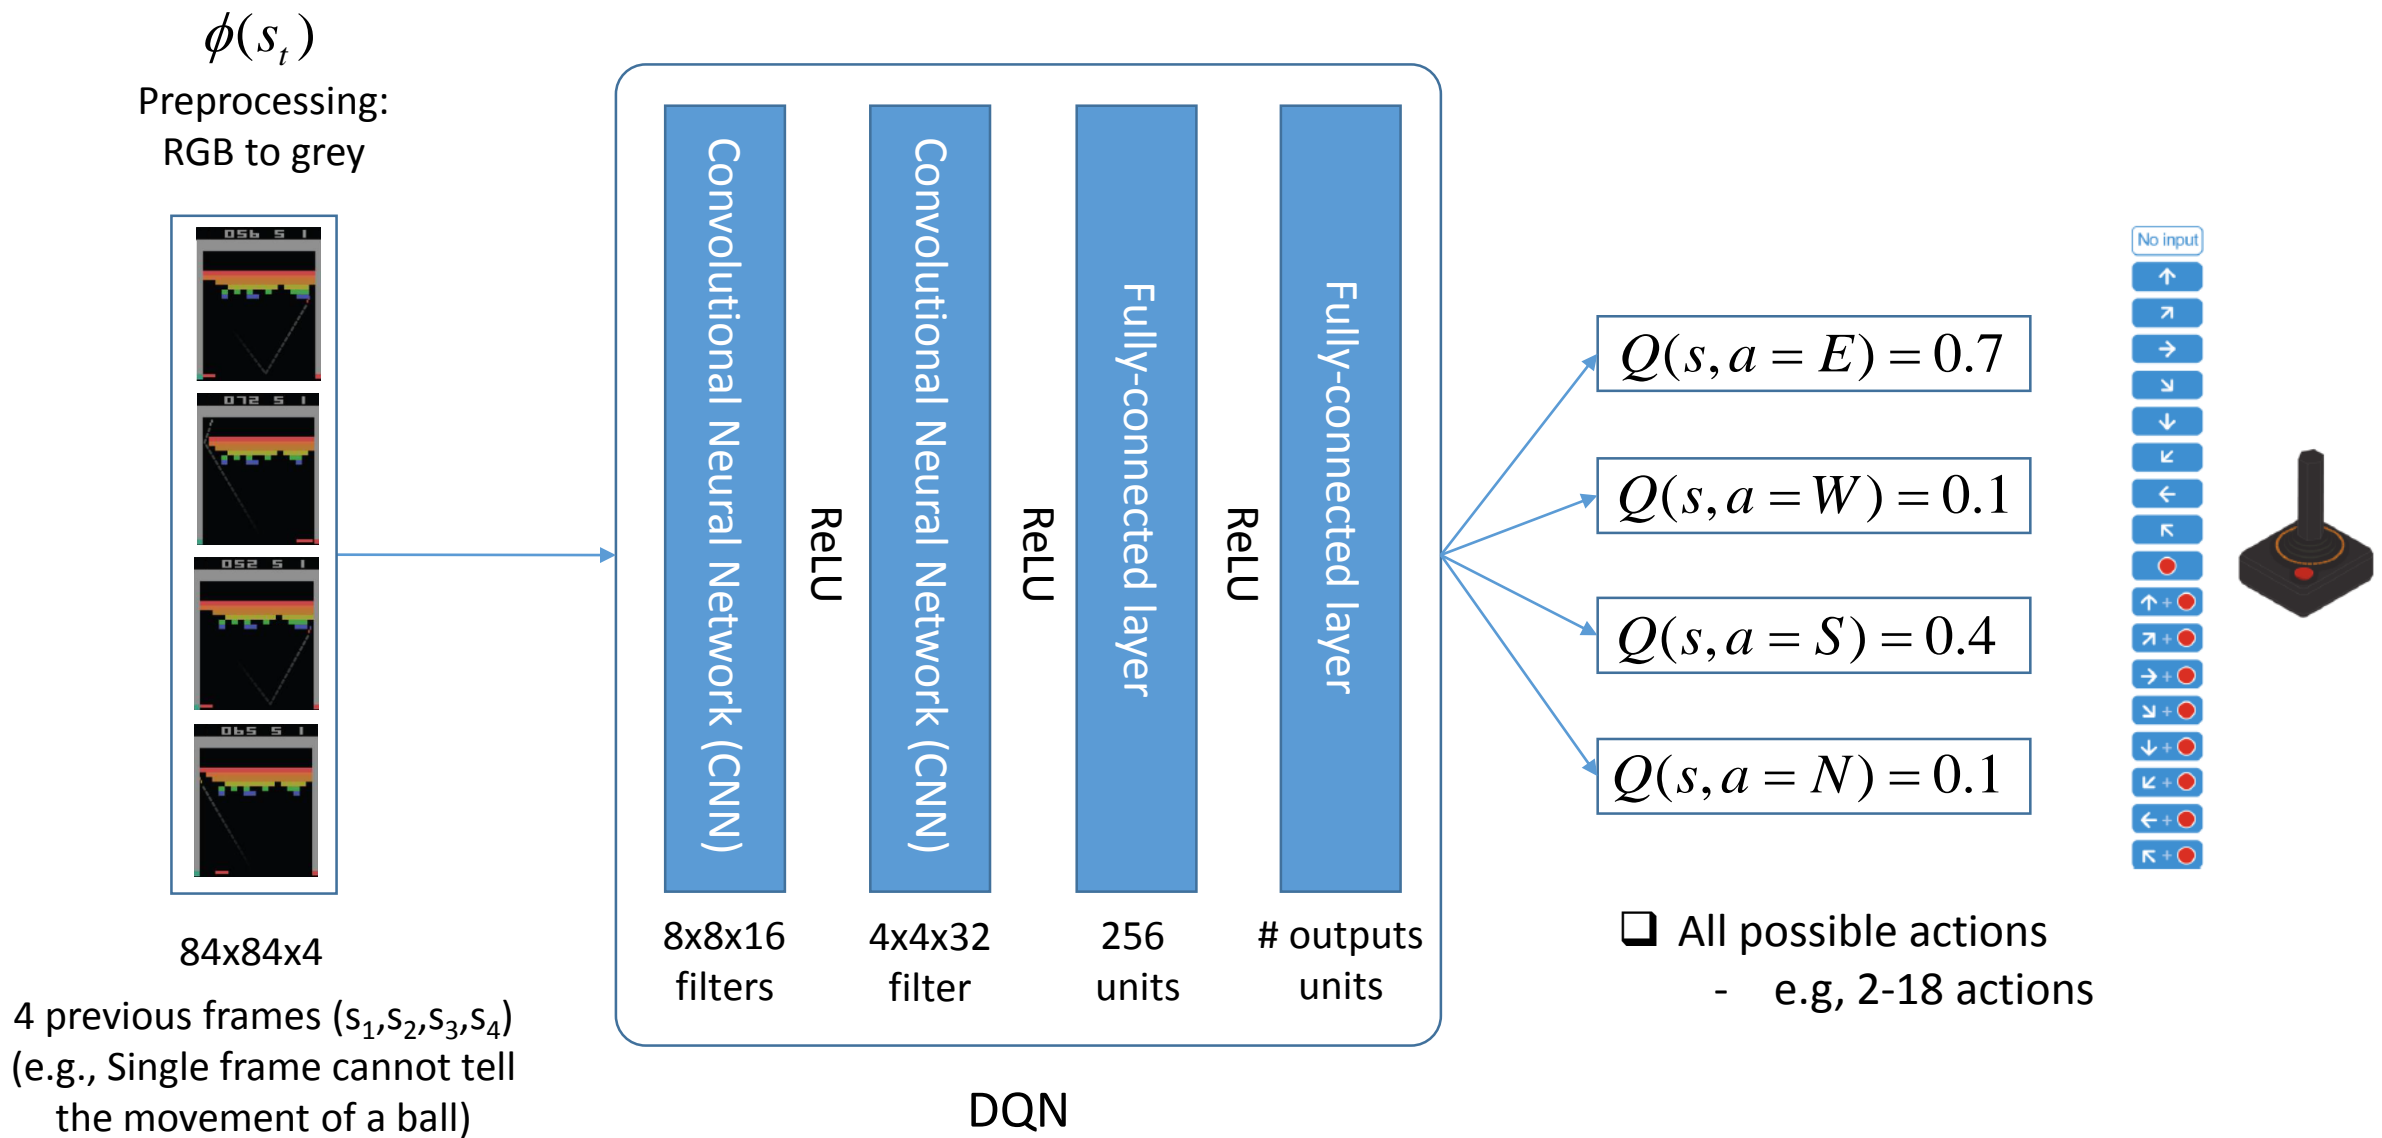
- ❑ In 2013, a team of DeepMind (AlphaGo) proposed convolutional neural networks as an approximation of the $Q(s, a)$ function.
- ❑ Then, it was named as **Deep Q-networks (DQN)**

$$\hat{Q}(s, a; \theta) \approx Q_{\pi}(s, a)$$

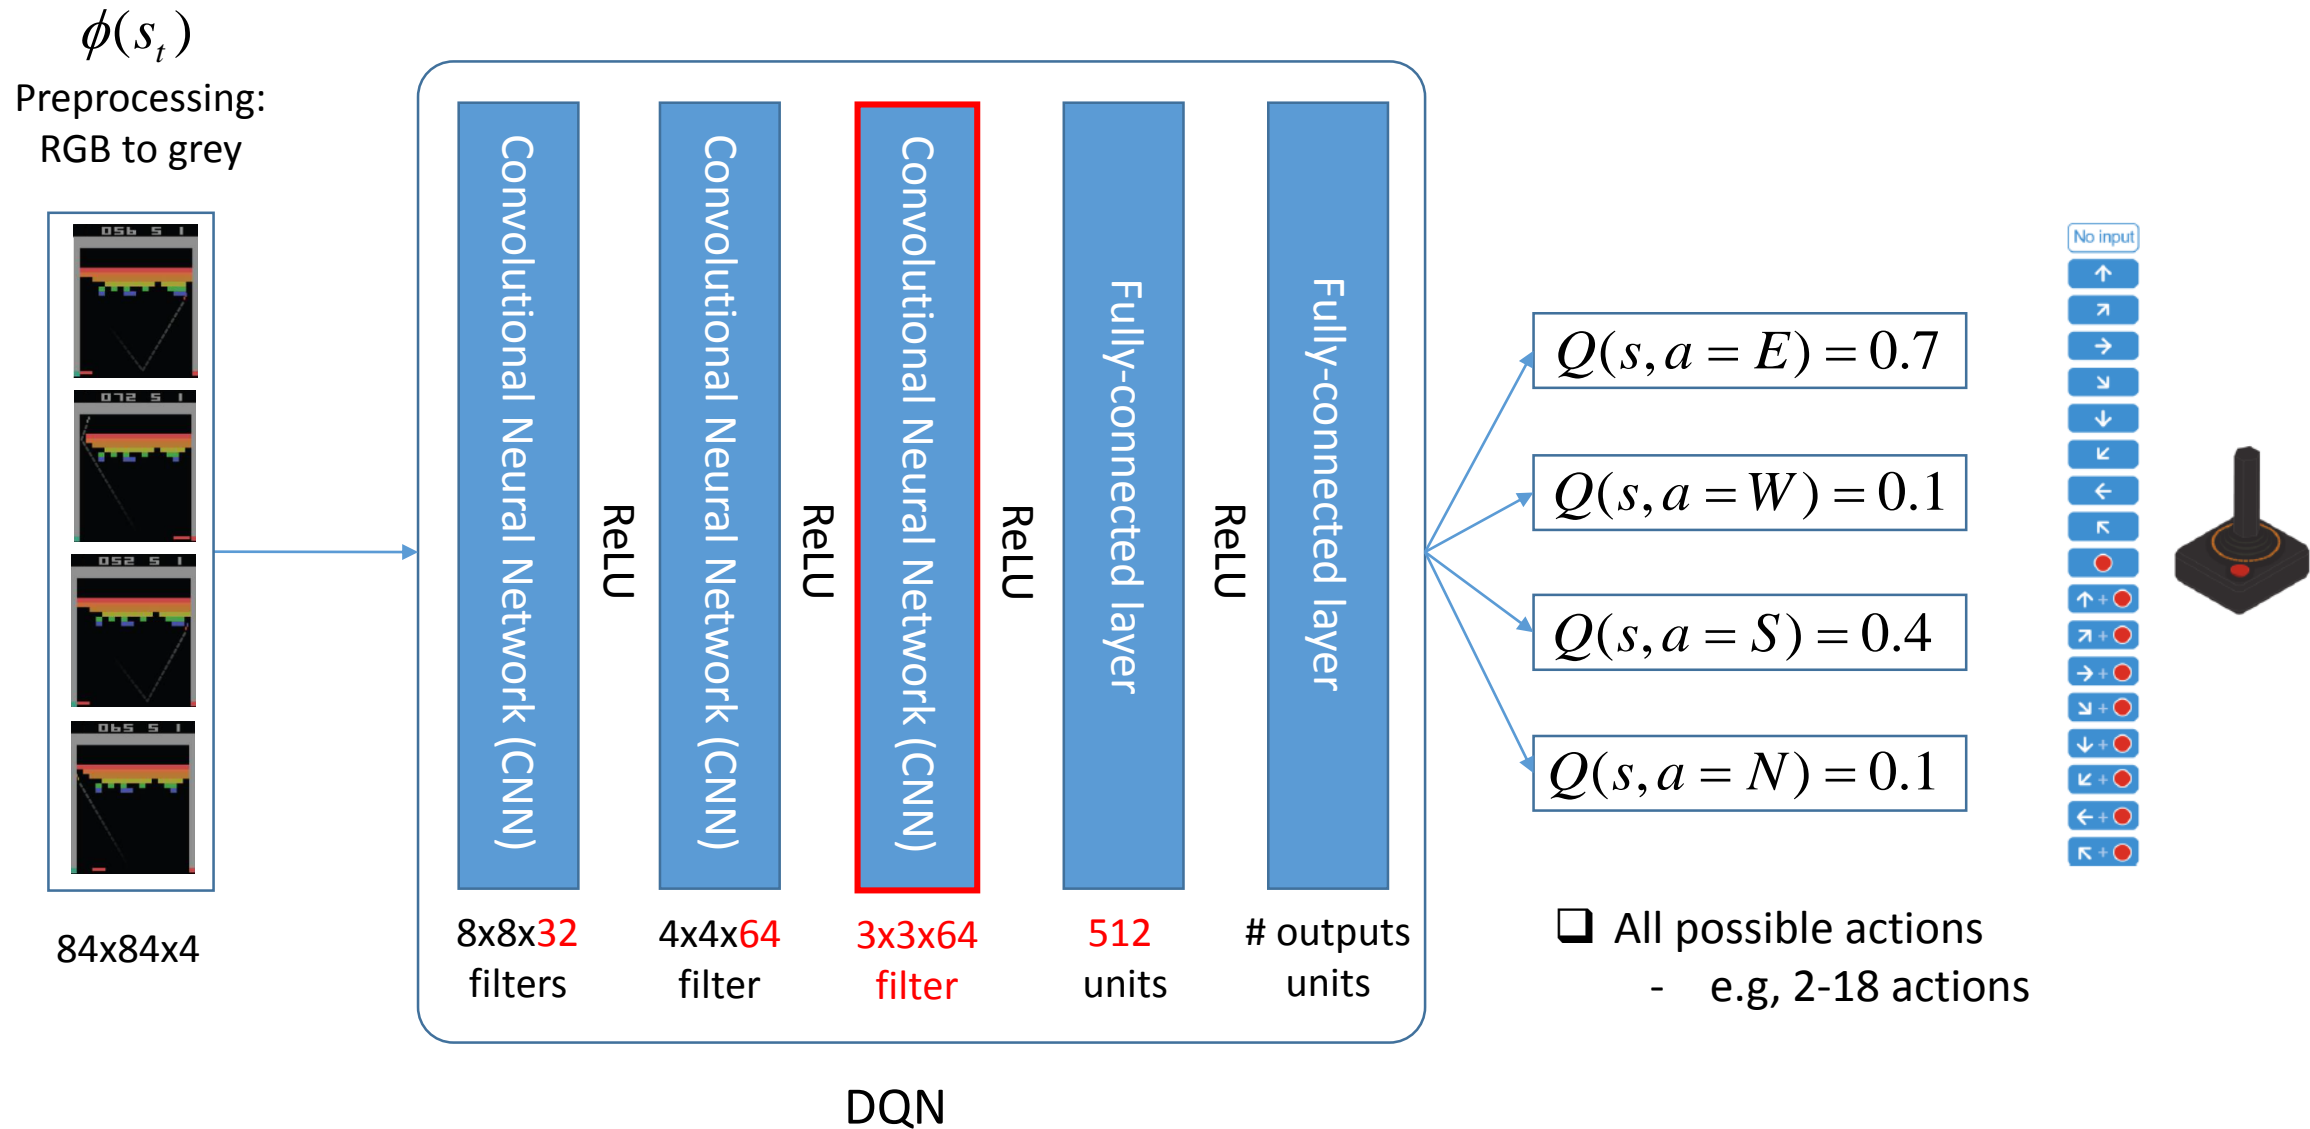
Hyper parameters + neural network parameters



Deep Q-Network architecture (2013)



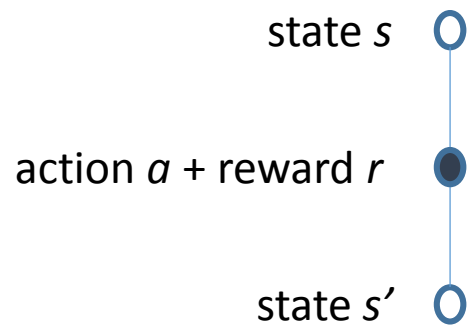
Deep Q-Network architecture (2015)



Loss function of the DQN

$$L_i(\theta_i) = \mathbb{E}_{(s,a,r,s') \sim U(D)} \left[\left(\underbrace{r + \gamma \max_{a'} Q(s', a' | \theta_i^-)}_{\text{Target future reward}} - \underbrace{Q(s, a, \theta_i)}_{\text{Expected future reward}} \right)^2 \right]$$

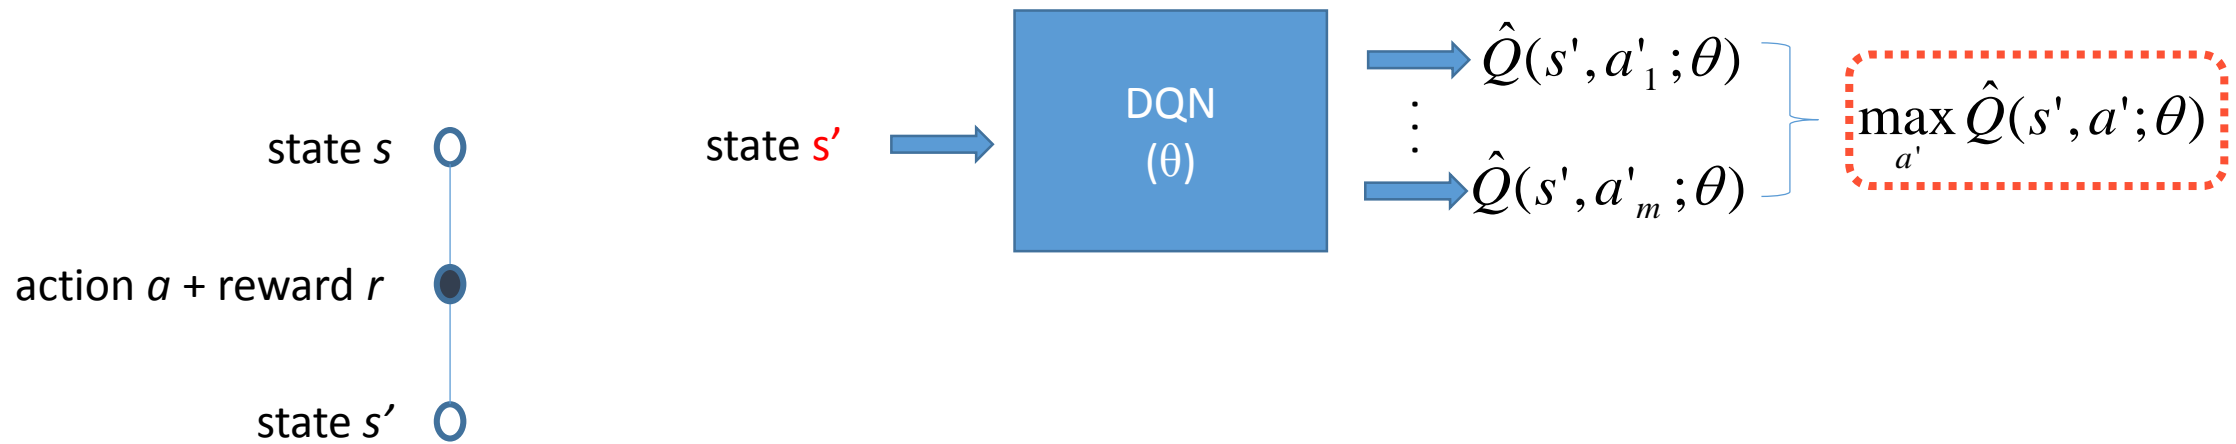
- Sample data (s,a,r,s') randomly drawn from data pool $U(D)$
- Experience Replay
- Two different neural networks
- Fixed Q-target



Loss function of the DQN

$$L_i(\theta_i) = \mathbb{E}_{(s,a,r,s') \sim U(D)} \left[\left(\underbrace{r + \gamma \max_{a'} Q(s', a' | \theta_i^-)}_{\text{Target future reward}} - \underbrace{Q(s, a, \theta_i)}_{\text{Expected future reward}} \right)^2 \right]$$

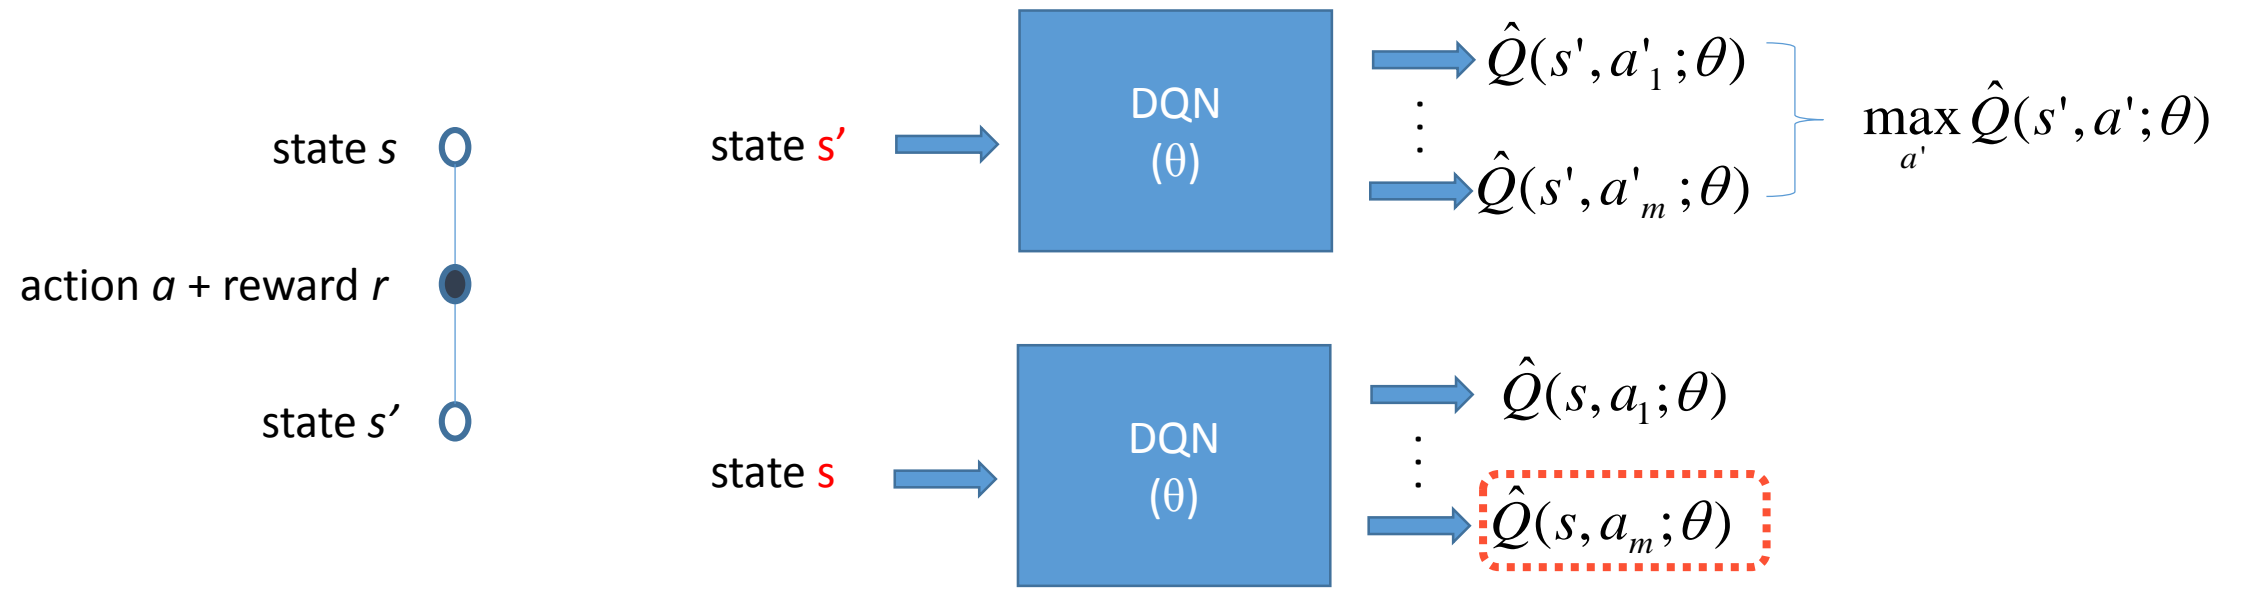
- Sample data (s,a,r,s') randomly drawn from data pool $U(D)$
- Experience Replay
- Two different neural networks
- Fixed Q-target



Loss function of the DQN

$$L_i(\theta_i) = \mathbb{E}_{(s,a,r,s') \sim U(D)} \left[\left(\underbrace{r + \gamma \max_{a'} Q(s', a' | \theta_i^-)}_{\text{Target future reward}} - \underbrace{Q(s, a, \theta_i)}_{\text{Expected future reward}} \right)^2 \right]$$

- Sample data (s,a,r,s') randomly drawn from data pool $U(D)$
- Experience Replay
- Two different neural networks
- Fixed Q-target

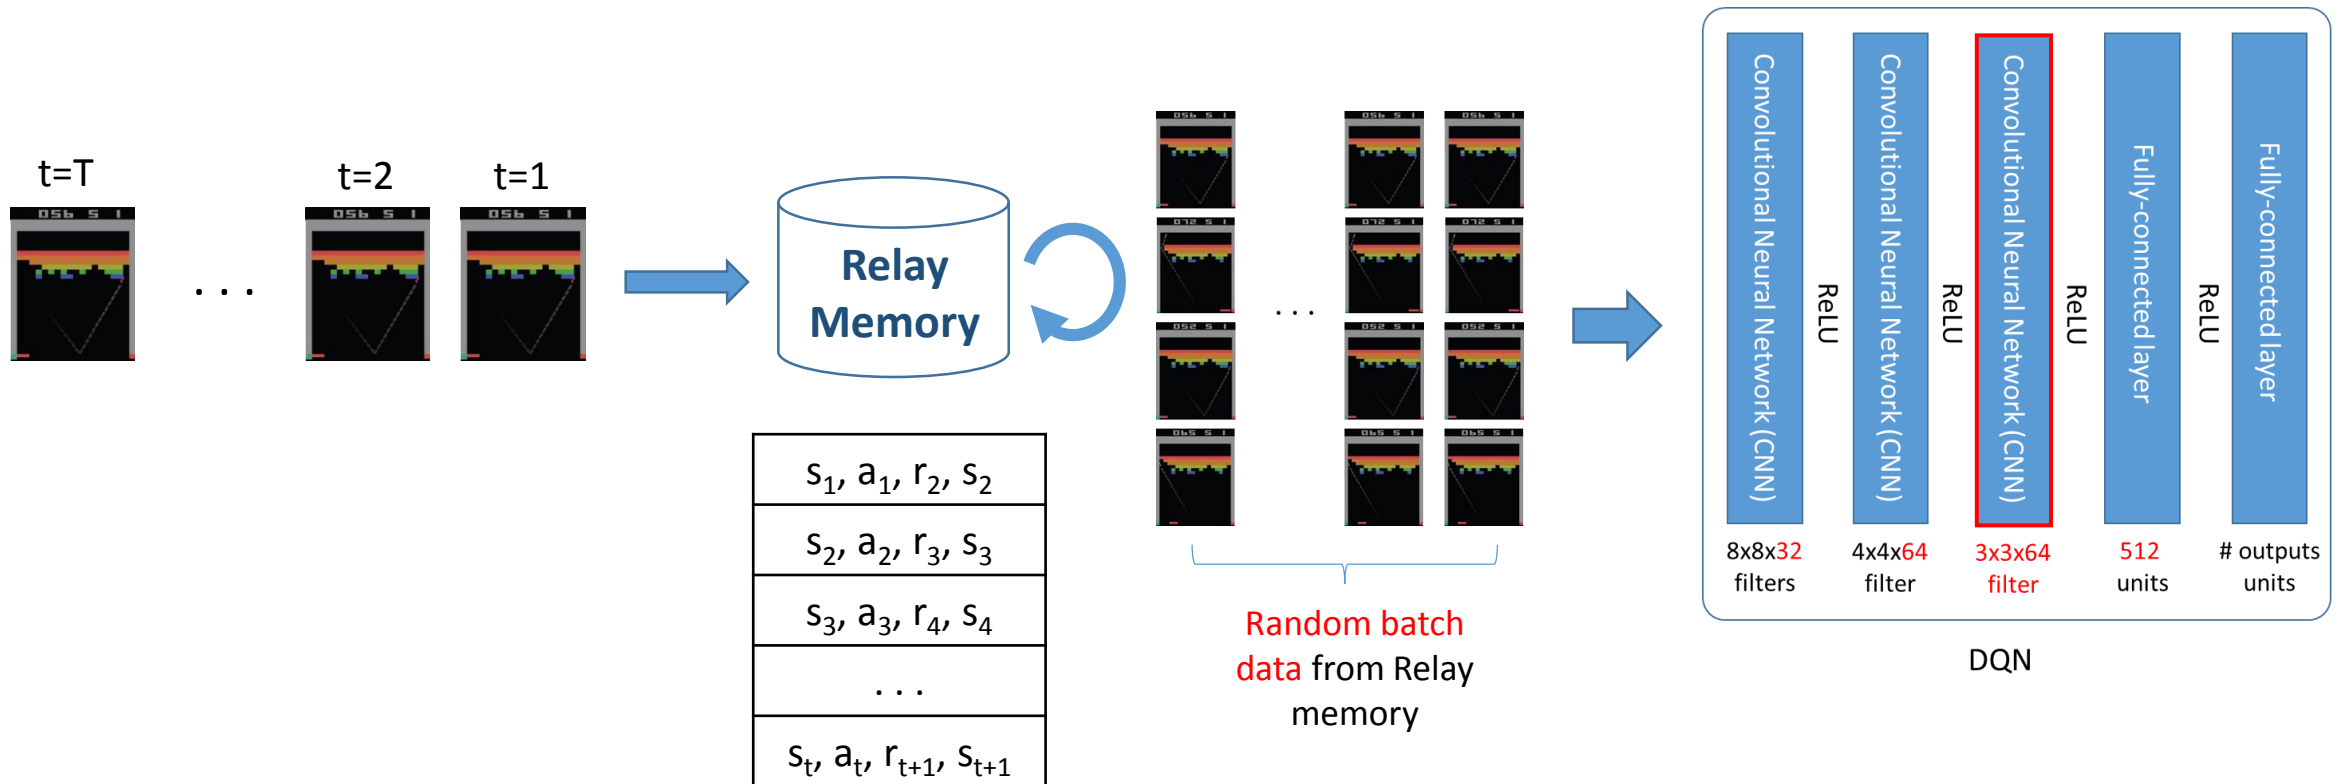


Deep Q-Network architecture (2013 & 2015)

- ❑ Neural networks were used previously for RL
 - Temporal Difference Learning and TD-Gammon (1992)
 - Deep Auto-Encoder Neural Networks in RL (2010)
- ❑ However, they were not successful due to oscillates or divergence of neural nets
- ❑ How does DQN handle this problem?
 - 1) Experience replay
 - 2) Fixed Q-targets
 - 3) Go deep

1) Experience replay

- ❑ Consecutive data frames are highly correlated
- ❑ Experience replay aims to remove the correlation between data samples

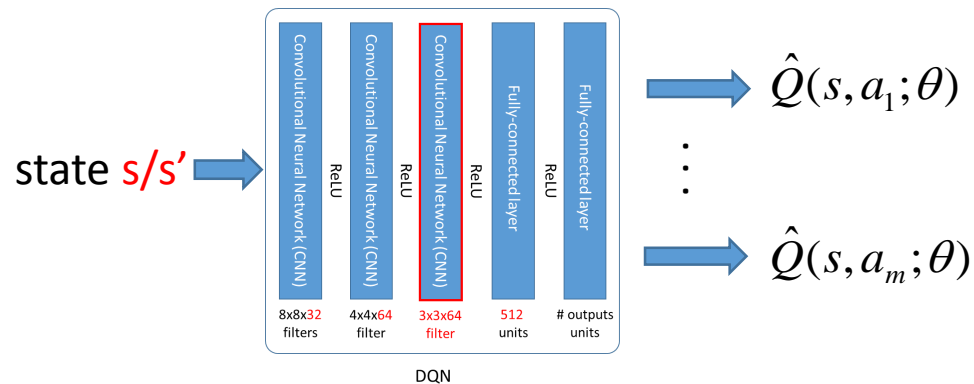


2) Fixed Q-target

- Originally the target future reward and the expected future reward are sharing the same neural net.

$$L_i(\theta_i) = E_{(s,a,r,s') \sim U(D)} \left[\left(r + \gamma \max_{a'} Q(s', a' | \theta_i) - Q(s, a, | \theta_i) \right)^2 \right]$$

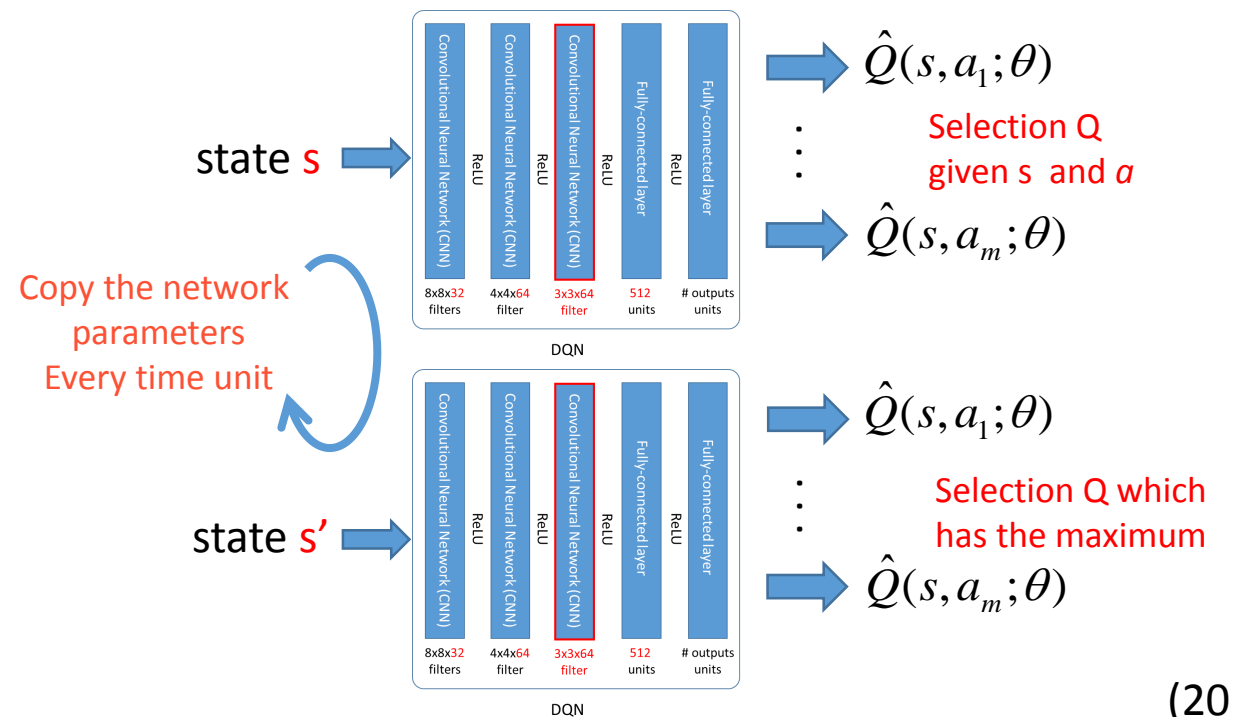
Same neural networks



(2013)

$$L_i(\theta_i) = E_{(s,a,r,s') \sim U(D)} \left[\left(r + \gamma \max_{a'} Q(s', a' | \theta_i^-) - Q(s, a, | \theta_i) \right)^2 \right]$$

Two different neural networks



(2015)

How effect they are?

2015 DQN

| Game | With replay, with target Q | With replay, without target Q | Without replay, with target Q | Without replay, without target Q |
|----------------|---------------------------------------|--|--|---|
| Breakout | 316.8 | 240.7 | 10.2 | 3.2 |
| Enduro | 1006.3 | 831.4 | 141.9 | 29.1 |
| River Raid | 7446.6 | 4102.8 | 2867.7 | 1453.0 |
| Seaquest | 2894.4 | 822.6 | 1003.0 | 275.8 |
| Space Invaders | 1088.9 | 826.3 | 373.2 | 302.0 |

Algorithm 1: deep Q-learning with experience replay.

Initialize replay memory D to capacity N -----> Data pool size and initialization

Initialize action-value function Q with random weights θ

Initialize target action-value function \hat{Q} with weights $\theta^- = \theta$

For episode = 1, M **do**

Initialize sequence $s_1 = \{x_1\}$ and preprocessed sequence $\phi_1 = \phi(s_1)$

For $t = 1, T$ **do**

With probability ε select a random action a_t

otherwise select $a_t = \operatorname{argmax}_a Q(\phi(s_t), a; \theta)$

Execute action a_t in emulator and observe reward r_t and image x_{t+1}

Set $s_{t+1} = s_t, a_t, x_{t+1}$ and preprocess $\phi_{t+1} = \phi(s_{t+1})$

Store transition $(\phi_t, a_t, r_t, \phi_{t+1})$ in D

Sample random minibatch of transitions $(\phi_j, a_j, r_j, \phi_{j+1})$ from D

Set $y_j = \begin{cases} r_j & \text{if episode terminates at step } j+1 \\ r_j + \gamma \max_{a'} \hat{Q}(\phi_{j+1}, a'; \theta^-) & \text{otherwise} \end{cases}$

Perform a gradient descent step on $(y_j - Q(\phi_j, a_j; \theta))^2$ with respect to the network parameters θ

Every C steps reset $\hat{Q} = Q$

End For

End For

Algorithm 1: deep Q-learning with experience replay.

Initialize replay memory D to capacity N -----> Data pool size and initialization
Initialize action-value function Q with random weights θ -----> Weight of 1st NN initialization
Initialize target action-value function \hat{Q} with weights $\theta^- = \theta$ -----> Weight of 2nd NN initialization

For episode = 1, M **do**

Initialize sequence $s_1 = \{x_1\}$ and preprocessed sequence $\phi_1 = \phi(s_1)$

For $t = 1, T$ **do**

With probability ε select a random action a_t

otherwise select $a_t = \operatorname{argmax}_a Q(\phi(s_t), a; \theta)$

Execute action a_t in emulator and observe reward r_t and image x_{t+1}

Set $s_{t+1} = s_t, a_t, x_{t+1}$ and preprocess $\phi_{t+1} = \phi(s_{t+1})$

Store transition $(\phi_t, a_t, r_t, \phi_{t+1})$ in D

Sample random minibatch of transitions $(\phi_j, a_j, r_j, \phi_{j+1})$ from D

Set $y_j = \begin{cases} r_j & \text{if episode terminates at step } j+1 \\ r_j + \gamma \max_{a'} \hat{Q}(\phi_{j+1}, a'; \theta^-) & \text{otherwise} \end{cases}$

Perform a gradient descent step on $(y_j - Q(\phi_j, a_j; \theta))^2$ with respect to the network parameters θ

Every C steps reset $\hat{Q} = Q$

End For

End For

Algorithm (DQN 2015)

Algorithm 1: deep Q-learning with experience replay.

Initialize replay memory D to capacity N -----> Data pool size and initialization

Initialize action-value function Q with random weights θ -----> Weight of 1st NN initialization

Initialize target action-value function \hat{Q} with weights $\theta^- = \theta$ -----> Weight of 2nd NN initialization

For episode = 1, M **do**

 Initialize sequence $s_1 = \{x_1\}$ and preprocessed sequence $\phi_1 = \phi(s_1)$ -----> Preprocessing, e.g., RGB to gray

For $t = 1, T$ **do**

 With probability ε select a random action a_t

 otherwise select $a_t = \operatorname{argmax}_a Q(\phi(s_t), a; \theta)$

 Execute action a_t in emulator and observe reward r_t and image x_{t+1}

 Set $s_{t+1} = s_t, a_t, x_{t+1}$ and preprocess $\phi_{t+1} = \phi(s_{t+1})$

 Store transition $(\phi_t, a_t, r_t, \phi_{t+1})$ in D

 Sample random minibatch of transitions $(\phi_j, a_j, r_j, \phi_{j+1})$ from D

 Set $y_j = \begin{cases} r_j & \text{if episode terminates at step } j+1 \\ r_j + \gamma \max_{a'} \hat{Q}(\phi_{j+1}, a'; \theta^-) & \text{otherwise} \end{cases}$

 Perform a gradient descent step on $(y_j - Q(\phi_j, a_j; \theta))^2$ with respect to the network parameters θ

 Every C steps reset $\hat{Q} = Q$

End For

End For

Algorithm (DQN 2015)

Algorithm 1: deep Q-learning with experience replay.

Initialize replay memory D to capacity N -----> Data pool size and initialization

Initialize action-value function Q with random weights θ -----> Weight of 1st NN initialization

Initialize target action-value function \hat{Q} with weights $\theta^- = \theta$ -----> Weight of 2nd NN initialization

For episode = 1, M **do**

 Initialize sequence $s_1 = \{x_1\}$ and preprocessed sequence $\phi_1 = \phi(s_1)$ -----> Preprocessing, e.g., RGB to gray

For $t = 1, T$ **do**

 With probability ε select a random action a_t } -----> Action selection using E-greedy: off-policy

 otherwise select $a_t = \operatorname{argmax}_a Q(\phi(s_t), a; \theta)$ }

 Execute action a_t in emulator and observe reward r_t and image x_{t+1}

 Set $s_{t+1} = s_t, a_t, x_{t+1}$ and preprocess $\phi_{t+1} = \phi(s_{t+1})$

 Store transition $(\phi_t, a_t, r_t, \phi_{t+1})$ in D

 Sample random minibatch of transitions $(\phi_j, a_j, r_j, \phi_{j+1})$ from D

 Set $y_j = \begin{cases} r_j & \text{if episode terminates at step } j+1 \\ r_j + \gamma \max_{a'} \hat{Q}(\phi_{j+1}, a'; \theta^-) & \text{otherwise} \end{cases}$

 Perform a gradient descent step on $(y_j - Q(\phi_j, a_j; \theta))^2$ with respect to the network parameters θ

 Every C steps reset $\hat{Q} = Q$

End For

End For

Algorithm (DQN 2015)

Algorithm 1: deep Q-learning with experience replay.

Initialize replay memory D to capacity N -----> Data pool size and initialization

Initialize action-value function Q with random weights θ -----> Weight of 1st NN initialization

Initialize target action-value function \hat{Q} with weights $\theta^- = \theta$ -----> Weight of 2nd NN initialization

For episode = 1, M **do**

 Initialize sequence $s_1 = \{x_1\}$ and preprocessed sequence $\phi_1 = \phi(s_1)$ -----> Preprocessing, e.g., RGB to gray

For $t = 1, T$ **do**

 With probability ε select a random action a_t } -----> Action selection using E-greedy: off-policy

 otherwise select $a_t = \operatorname{argmax}_a Q(\phi(s_t), a; \theta)$ }

 Execute action a_t in emulator and observe reward r_t and image x_{t+1}

 Set $s_{t+1} = s_t, a_t, x_{t+1}$ and preprocess $\phi_{t+1} = \phi(s_{t+1})$

 Store transition $(\phi_t, a_t, r_t, \phi_{t+1})$ in D

 Sample random minibatch of transitions $(\phi_j, a_j, r_j, \phi_{j+1})$ from D -----> Experience replay

 Set $y_j = \begin{cases} r_j & \text{if episode terminates at step } j+1 \\ r_j + \gamma \max_{a'} \hat{Q}(\phi_{j+1}, a'; \theta^-) & \text{otherwise} \end{cases}$

 Perform a gradient descent step on $(y_j - Q(\phi_j, a_j; \theta))^2$ with respect to the network parameters θ

 Every C steps reset $\hat{Q} = Q$

End For

End For

Algorithm (DQN 2015)

Algorithm 1: deep Q-learning with experience replay.

Initialize replay memory D to capacity N -----> Data pool size and initialization

Initialize action-value function Q with random weights θ -----> Weight of 1st NN initialization

Initialize target action-value function \hat{Q} with weights $\theta^- = \theta$ -----> Weight of 2nd NN initialization

For episode = 1, M **do**

 Initialize sequence $s_1 = \{x_1\}$ and preprocessed sequence $\phi_1 = \phi(s_1)$ -----> Preprocessing, e.g., RGB to gray

For $t = 1, T$ **do**

 With probability ε select a random action a_t } -----> Action selection using E-greedy: off-policy

 otherwise select $a_t = \operatorname{argmax}_a Q(\phi(s_t), a; \theta)$ }

 Execute action a_t in emulator and observe reward r_t and image x_{t+1}

 Set $s_{t+1} = s_t, a_t, x_{t+1}$ and preprocess $\phi_{t+1} = \phi(s_{t+1})$

 Store transition $(\phi_t, a_t, r_t, \phi_{t+1})$ in D

 Sample random minibatch of transitions $(\phi_j, a_j, r_j, \phi_{j+1})$ from D -----> Experience replay

 Set $y_j = \begin{cases} r_j & \text{if episode terminates at step } j+1 \\ r_j + \gamma \max_{a'} \hat{Q}(\phi_{j+1}, a'; \theta^-) & \text{otherwise} \end{cases}$ -----> Target future reward is obtained from NN (θ^-)

 Perform a gradient descent step on $(y_j - Q(\phi_j, a_j; \theta))^2$ with respect to the network parameters θ

 Every C steps reset $\hat{Q} = Q$

End For

End For

Algorithm (DQN 2015)

Algorithm 1: deep Q-learning with experience replay.

Initialize replay memory D to capacity N -----> Data pool size and initialization

Initialize action-value function Q with random weights θ -----> Weight of 1st NN initialization

Initialize target action-value function \hat{Q} with weights $\theta^- = \theta$ -----> Weight of 2nd NN initialization

For episode = 1, M **do**

 Initialize sequence $s_1 = \{x_1\}$ and preprocessed sequence $\phi_1 = \phi(s_1)$ -----> Preprocessing, e.g., RGB to gray

For $t = 1, T$ **do**

 With probability ε select a random action a_t } -----> Action selection using E-greedy: off-policy

 otherwise select $a_t = \operatorname{argmax}_a Q(\phi(s_t), a; \theta)$ }

 Execute action a_t in emulator and observe reward r_t and image x_{t+1}

 Set $s_{t+1} = s_t, a_t, x_{t+1}$ and preprocess $\phi_{t+1} = \phi(s_{t+1})$

 Store transition $(\phi_t, a_t, r_t, \phi_{t+1})$ in D

 Sample random minibatch of transitions $(\phi_j, a_j, r_j, \phi_{j+1})$ from D -----> Experience replay

 Set $y_j = \begin{cases} r_j & \text{if episode terminates at step } j+1 \\ r_j + \gamma \max_{a'} \hat{Q}(\phi_{j+1}, a'; \theta^-) & \text{otherwise} \end{cases}$ -----> Target future reward is obtained from NN (θ^-)

 Perform a gradient descent step on $(y_j - Q(\phi_j, a_j; \theta))^2$ with respect to the -----> Update NN (θ) without changing NN (θ^-)

 network parameters θ

 Every C steps reset $\hat{Q} = Q$

End For

End For

Algorithm (DQN 2015)

Algorithm 1: deep Q-learning with experience replay.

Initialize replay memory D to capacity N -----> Data pool size and initialization

Initialize action-value function Q with random weights θ -----> Weight of 1st NN initialization

Initialize target action-value function \hat{Q} with weights $\theta^- = \theta$ -----> Weight of 2nd NN initialization

For episode = 1, M **do**

 Initialize sequence $s_1 = \{x_1\}$ and preprocessed sequence $\phi_1 = \phi(s_1)$ -----> Preprocessing, e.g., RGB to gray

For $t = 1, T$ **do**

 With probability ε select a random action a_t } -----> Action selection using E-greedy: off-policy

 otherwise select $a_t = \operatorname{argmax}_a Q(\phi(s_t), a; \theta)$ }

 Execute action a_t in emulator and observe reward r_t and image x_{t+1}

 Set $s_{t+1} = s_t, a_t, x_{t+1}$ and preprocess $\phi_{t+1} = \phi(s_{t+1})$

 Store transition $(\phi_t, a_t, r_t, \phi_{t+1})$ in D

 Sample random minibatch of transitions $(\phi_j, a_j, r_j, \phi_{j+1})$ from D -----> Experience replay

 Set $y_j = \begin{cases} r_j & \text{if episode terminates at step } j+1 \\ r_j + \gamma \max_{a'} \hat{Q}(\phi_{j+1}, a'; \theta^-) & \text{otherwise} \end{cases}$ -----> Target future reward is obtained from NN (θ^-)

 Perform a gradient descent step on $(y_j - Q(\phi_j, a_j; \theta))^2$ with respect to the network parameters θ -----> Update NN (θ) without changing NN (θ^-)

 Every C steps reset $\hat{Q} = Q$ -----> Replace NN (θ^-) with NN (θ) every C steps

End For

End For

Policy Gradient (PG)

Difference between DQN and Policy Gradient

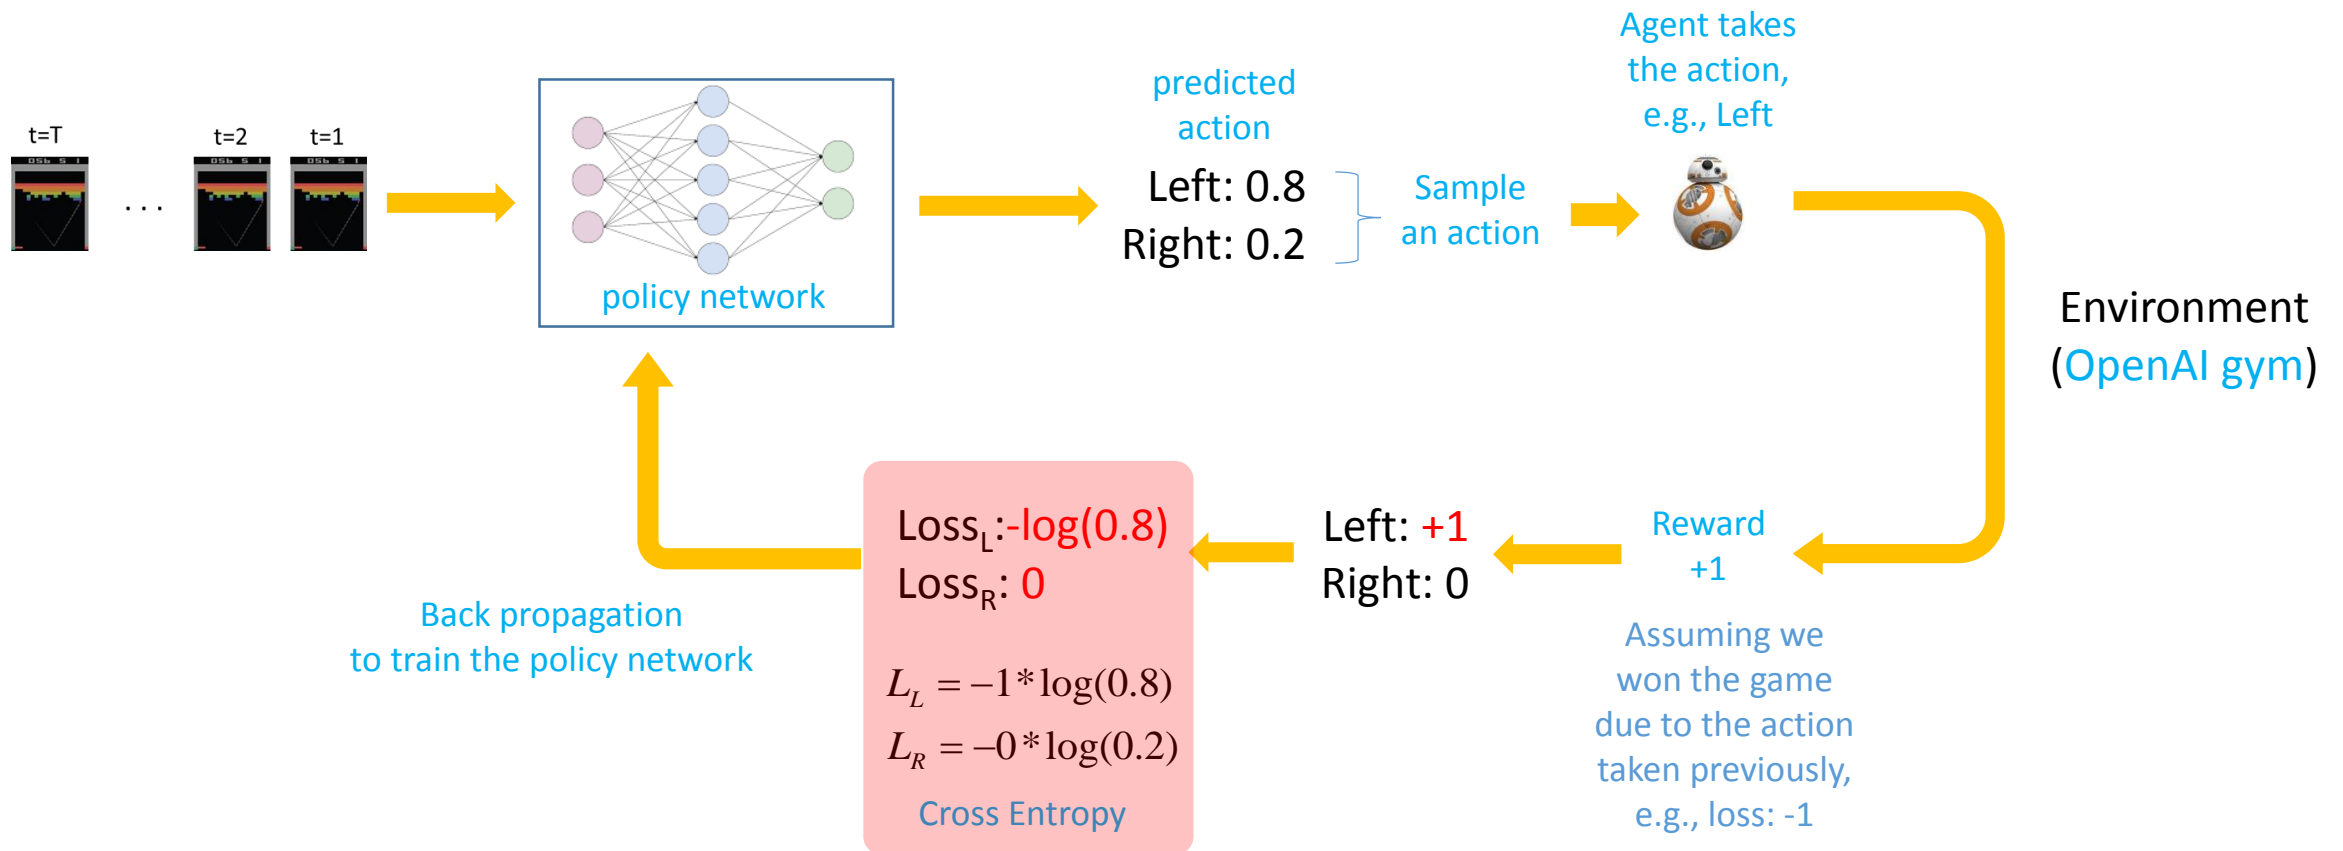
- When we can approximate Q function for all states and action pairs, we can obtain the optimal π^* by following way:

$$\pi^*(s) = \arg \max_a Q(s, a) \quad \text{: optimal policy}$$

- Policy Gradient (PG) directly optimizes the policy function π without obtaining Q function.
 - Similar to DQN, PG can also use a neural network (Policy Network): the output is the probability of each action at given state.

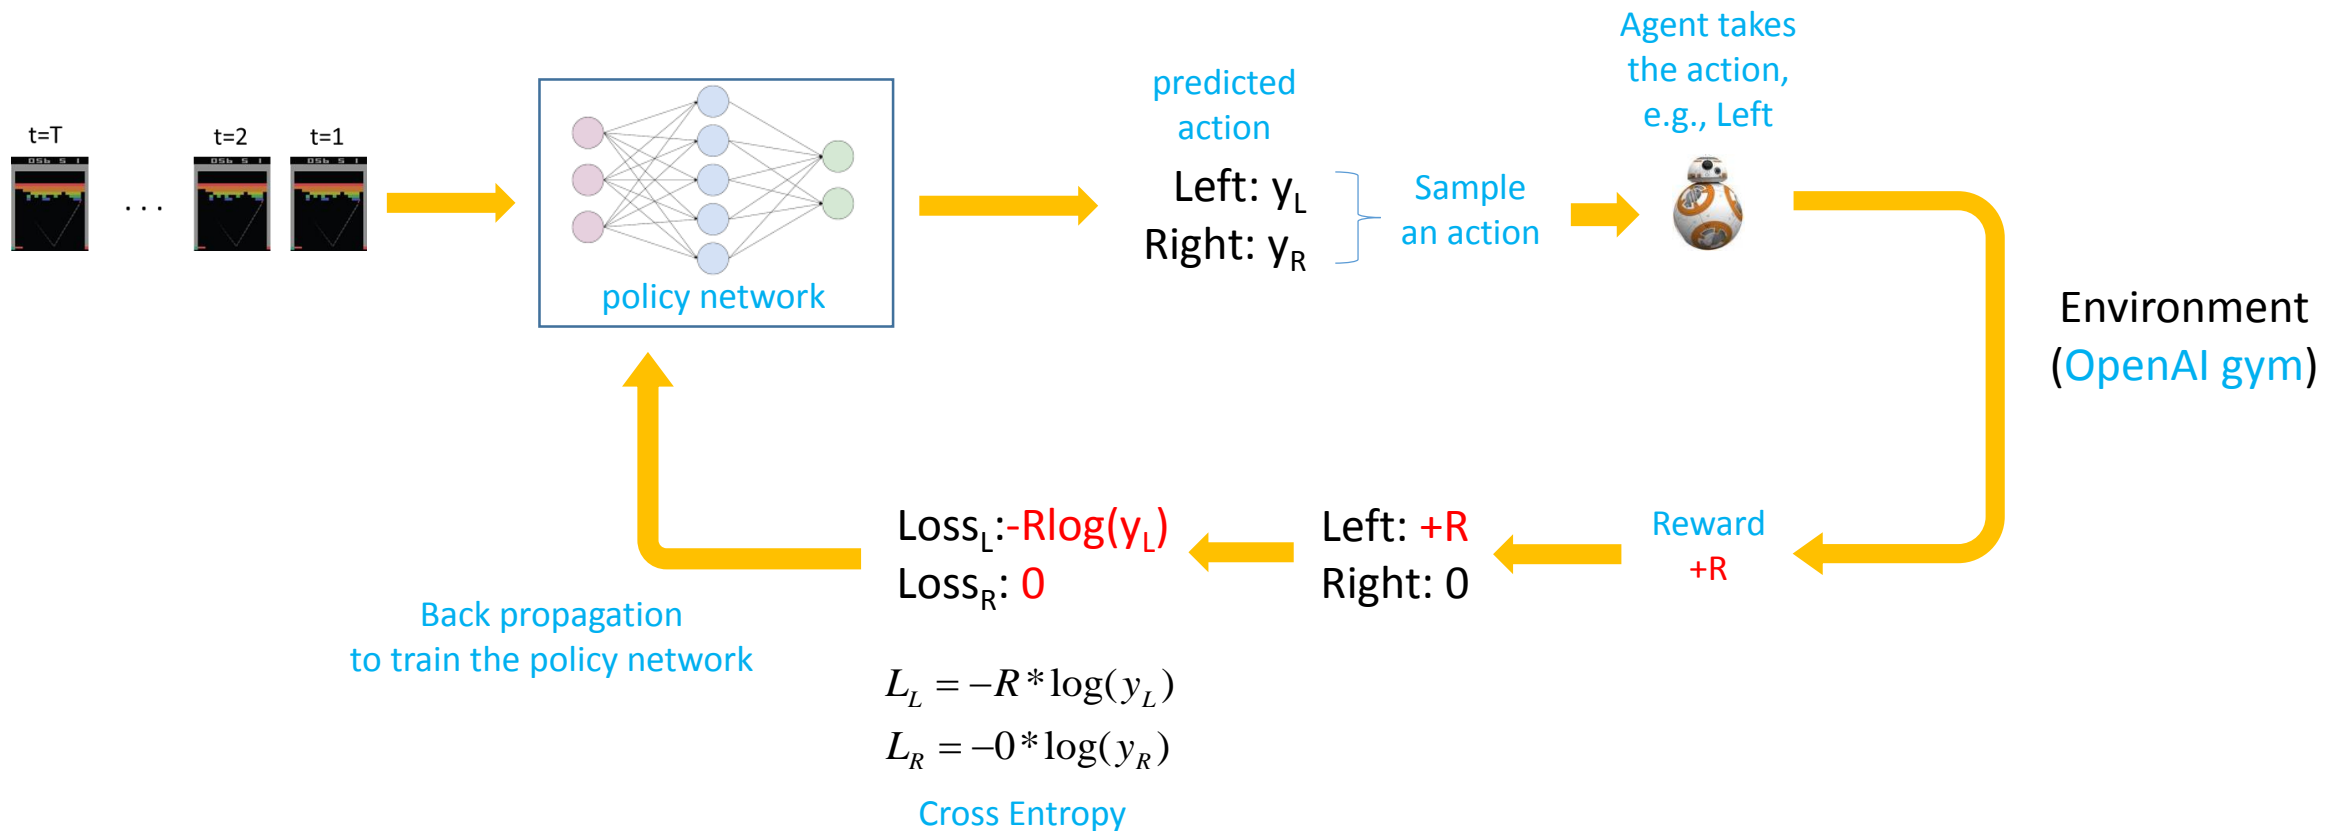
Loss function for Policy Network (PN): Cross Entropy

□ First, let's see the overall operation of policy gradient method



Loss function for Policy Network (PN): Derivation

- Let's generalize the loss function by introducing Reward.



Loss function for Policy Network (PN): Derivation - cont

- ❑ PG aims to obtain an optimum policy which maximizes the future reward.
- ❑ In the previous slide, we want to train PN in a way that
 - When an agent follows the policy given by the outcome of the PN, it expects high future reward.

$$L_i(\theta_i) = \log(\pi_\theta(s, a)) \cdot R$$

- Expected future reward triggered by the sampled action



- An action is sampled from a policy
- NN models the policy distribution

- “-” sign disappeared because we want to maximize the reward (good one has large reward)



- This equation says: the parameter θ of PN is updated by optimizing the policy π which maximizes future reward.

Q learning vs Policy gradient

| Q learning | Policy gradient |
|--|---|
| <ul style="list-style-type: none">Learning $Q(s,a)$: modeling (Reward) values of actions- Value based approach: learning Q values | <ul style="list-style-type: none">Learning $\pi(a)$: modeling probability of actions- Policy based approach: learning policy directly |
| <ul style="list-style-type: none">Deterministic policies:- e.g., cannot model rock-paper-scissors game | <ul style="list-style-type: none">Stochastic policies- e.g., can model rock-paper-scissors game |
| <ul style="list-style-type: none">Off-policy: an action is taken greedily- Greed search to calculate $Q(s,a)$ and then determine a policy | <ul style="list-style-type: none">On-policy: an action is taken with a policy- Following a trajectory created by a policy and update it with given reward at the end. |
| <ul style="list-style-type: none">Learning update occurred step-by-step (bootstrapping)- Low variance but high bias | <ul style="list-style-type: none">Learning update occurred episode-by-episode- High variance but low bias |



High bias




High variance

Action Critic (AC)

Action critic = Q learning + Policy gradient

- It aims to deal with following two problems in PG.
 - 1) PG uses episode-by-episode learning update, which disables on-line learning.
 - 2) PG tends to produce a policy with high variance.

Policy Gradient approach
Monte Carlo (MC) Learning
Episode-by-episode




1) Learning update step-by-step: Policy gradient theorem

- For the first problem, let's **change the reward function R to $Q(s, a)$ function** and so learning update can be done step-by-step, which **enables on-line learning**.
 - Proof in “*policy gradient methods for reinforcement learning with function approximation*”


Monte Carlos (MC) learning

$$L_i(\theta_i) = \log(\pi_\theta(s, a)) \cdot R$$



Temporal-Difference (TD) learning

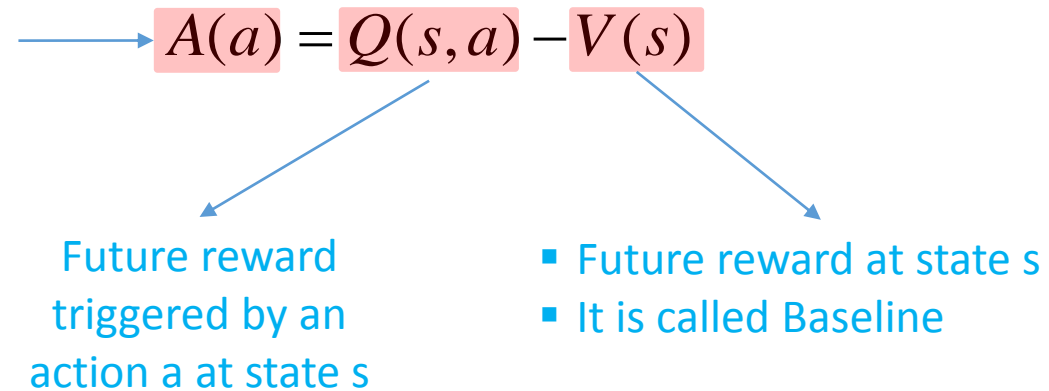
$$L_i(\theta_i) = \log(\pi_\theta(s, a)) \cdot Q(s, a)$$



2) Reducing variance

- ❑ For the second problem, let's introduce “**advantage**” which replaces “**Q function**”
- ❑ Well, it is a kind of normalization process if you see its definition below.

- How good is the action a at state s comparing to the average future reward at state s ?



$$L_i(\theta_i) = \log(\pi_\theta(s, a)) \cdot Q(s, a) \quad \longrightarrow \quad L_i(\theta_i) = \log(\pi_\theta(s, a)) \cdot A(a)$$
$$= \log(\pi_\theta(s, a)) \cdot (Q(s, a) - V(s))$$
$$= \log(\pi_\theta(s_t, a_t)) \cdot (r_{t+1} + \gamma V(s_{t+1}) - V(s_t))$$

- Actor-Critic is a policy gradient approach which updates a policy in each step
 - 1) Actor determines a policy
 - 2) Critic determines a value function for future reward

$$L_i(\theta_i) = \log(\pi_\theta(s_t, a_t)) \cdot (r_{t+1} + \gamma V(s_{t+1}) - V(s_t))$$


- We call it “Critic” because it *critically(?)* evaluates how good the action taken.
- $V(s)$ needs to be found to calculate this part
- A neural network can be used to approximate this value

- We call it “Actor” because it determines its action policies!
- Policy needs to be determined based on the result from Critic
- A neural network can be used to approximate the policy

A3C: Asynchronous Advantage Actor-Critic

Two problems in Actor Critic

- ❑ High bias due to every one step update
- ❑ Exploration issue
 - DQN uses e-greedy approach to handle exploration issue.
 - However, policy gradient approach; Actor Critic does not have the mechanism.
 - Stochastic behavior of a policy function can handle the exploration issue partially.



1) High Bias every one step update

- ❑ A3C introduces **multi step updates** to handle the problem
- ❑ There can be several variations!

$$L(\theta) = \log(\pi_{\theta}(s_t, a_t)) \cdot (r_{t+1} + \gamma V(s_{t+1}) - V(s_t))$$


$$L(\theta) = \log(\pi_{\theta}(s_t, a_t)) \cdot (r_{t+1} + r_{t+2} + \gamma V(s_{t+2}) - V(s_t))$$

$$L(\theta) = \log(\pi_{\theta}(s_t, a_t)) \cdot (r_{t+1} + r_{t+2} + r_{t+3} + \gamma V(s_{t+3}) - V(s_t))$$

$$L(\theta) = \log(\pi_{\theta}(s_t, a_t)) \cdot (r_{t+1} + r_{t+2} + r_{t+3} + \gamma V(s_{t+3}) - V(s_t))$$

$$+ \log(\pi_{\theta}(s_t, a_t)) \cdot (r_{t+2} + r_{t+3} + \gamma V(s_{t+3}) - V(s_t))$$

$$+ \log(\pi_{\theta}(s_t, a_t)) \cdot (r_{t+3} + \gamma V(s_{t+3}) - V(s_t))$$



2) Exploration issue

- A3C includes “**entropy of the policy π** ” to the loss function in order to improve exploration by discouraging premature convergence to suboptimal deterministic policies.

$$L(\theta) = \log(\pi_{\theta}(s_t, a_t)) \cdot (r_{t+1} + \gamma V(s_{t+1}) - V(s_t)) + \beta H \pi_{\theta}(s_t, a_t)$$

- This term defines the probability distribution of actions at each state

- Entropy regularization term
- This term tries to uniformize the probability distribution of actions defined in the first term.
 - Entropy is maximized when all actions from the policy π are same.
 - It aims to occur all action with equal probability (exploration)

Hand-on Experience


OPENAI GYM

❑ OpenAI :

- A non-profit artificial intelligence (AI) research company that aims to promote and develop friendly AI in such a way to benefit humanity as a whole.
- In October 2015, Alon Musk et al founded the organization.
- On April 2016, OpenAI released a public beta of “OpenAI Gym”, its platform for reinforcement learning research.


❑ OpenAI Gym

- A toolkit for developing and comparing reinforcement learning algorithms
- <https://github.com/openai/gym>
- <https://gym.openai.com/>




Gym

Gym is a toolkit for developing and comparing reinforcement learning algorithms. It supports teaching agents everything from walking to playing games like Pong or Pinball.





RL algorithm


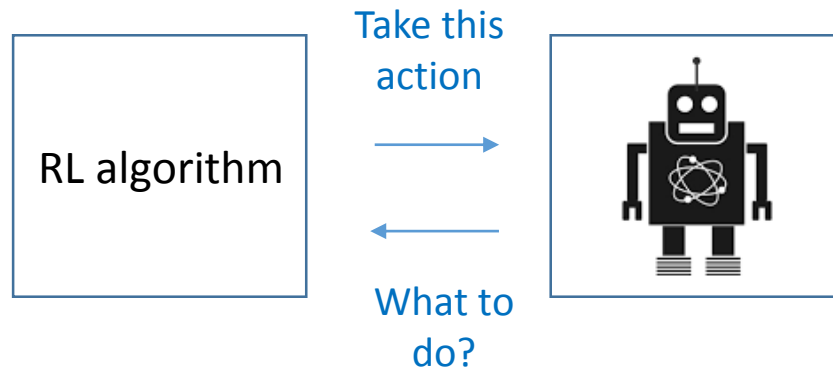



RL algorithm


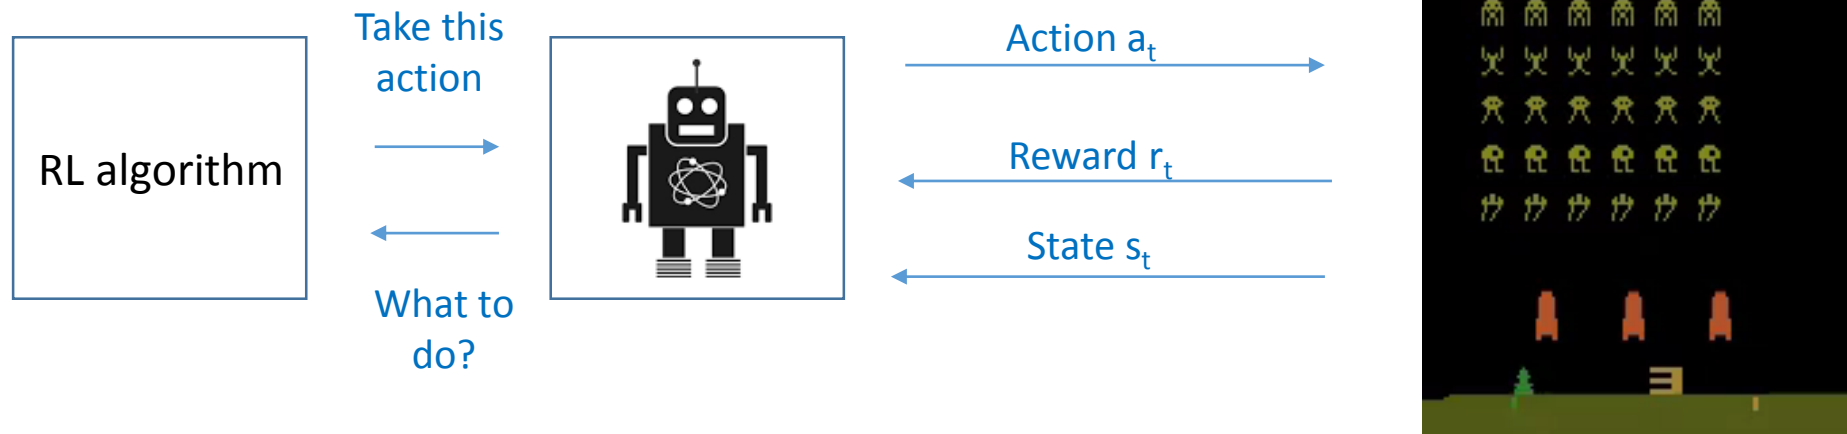


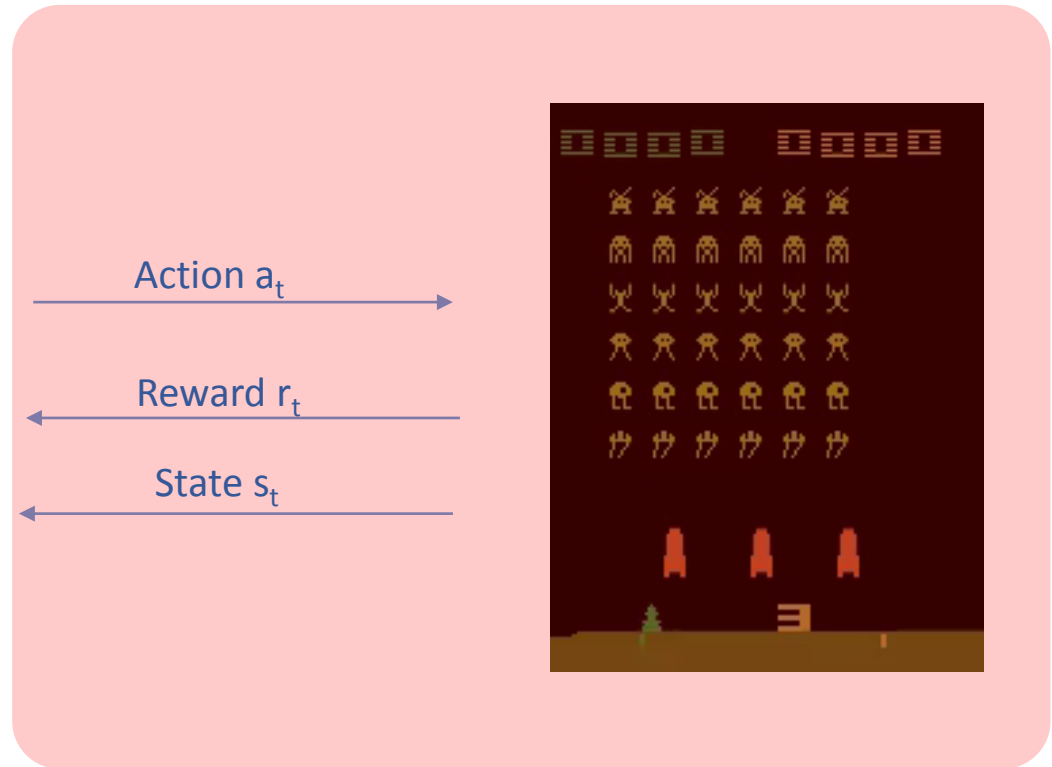
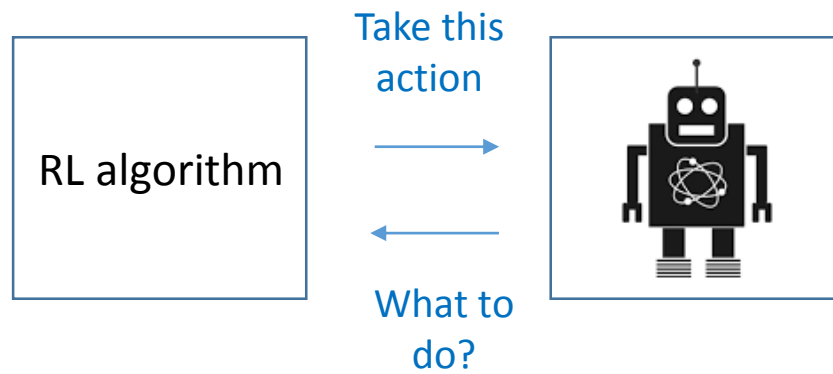
What to do?





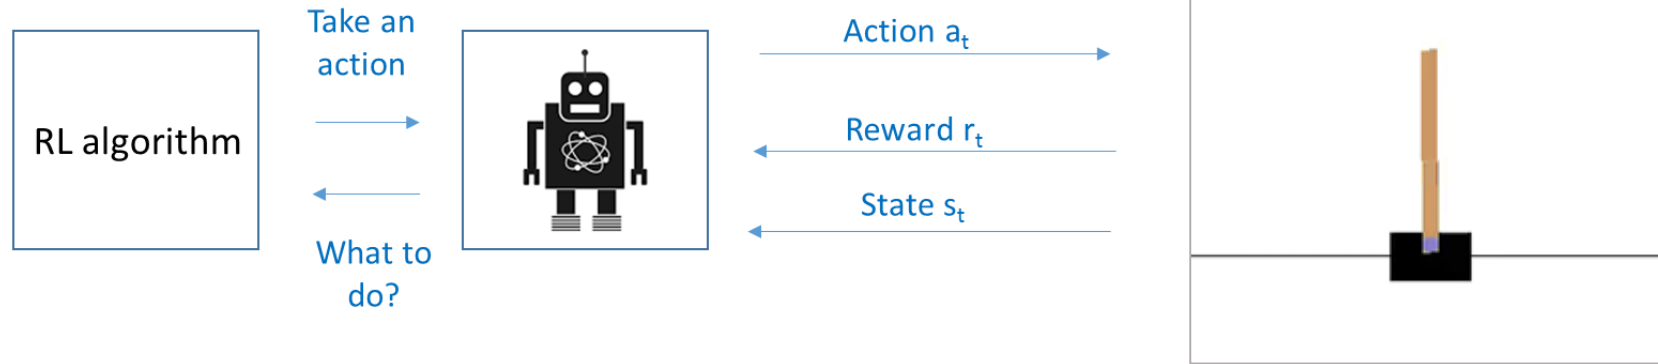




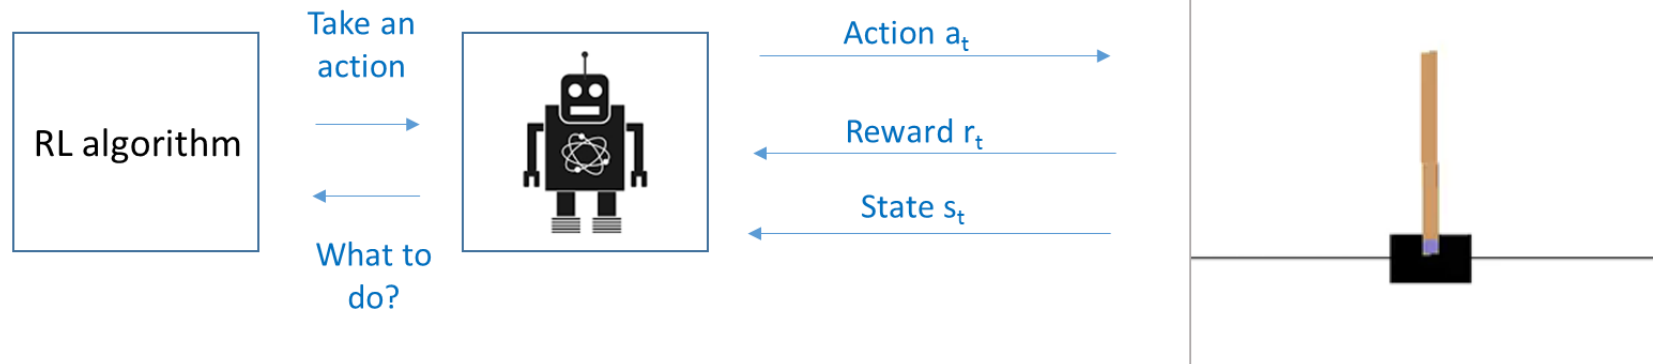


Gym framework

OpenAI Gym framework: CartPole game



OpenAI Gym framework: CartPole game



[-0.00907215 0.01075882 -0.0133813 0.01925749]



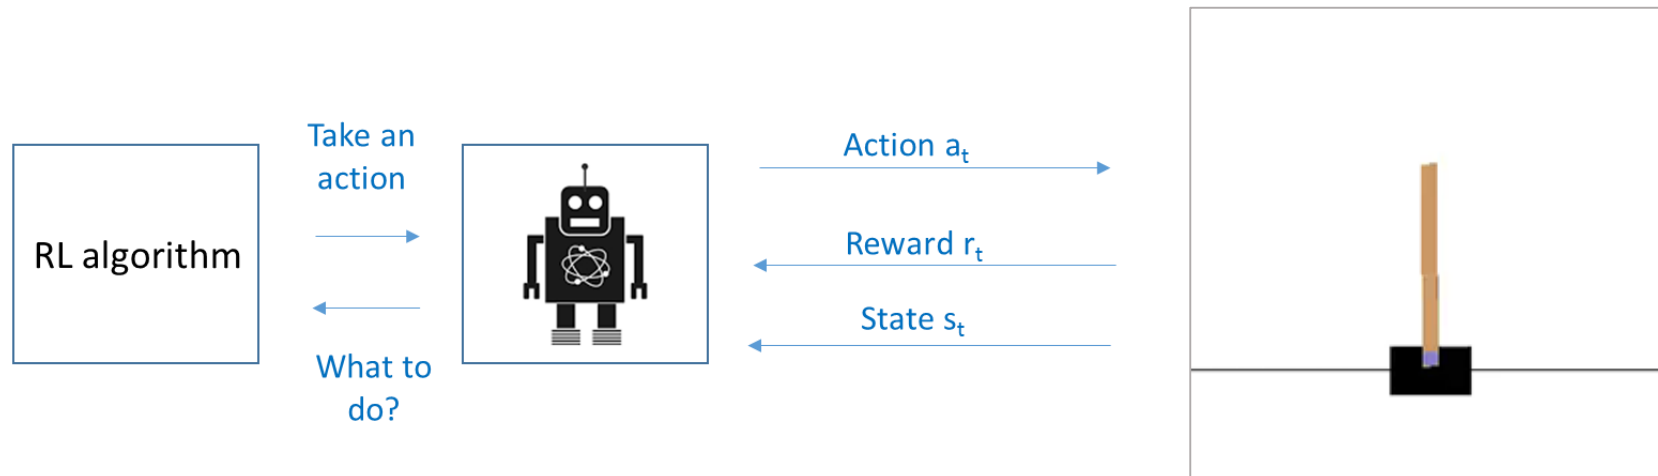
0:left, 1:right

```
import gym
env = gym.make('CartPole-v0')
for i_episode in range(20):
    observation = env.reset()
    for t in range(100):
        env.render()
        print(observation)
        action = env.action_space.sample()
        observation, reward, done, info = env.step(action)
        if done:
            print("Episode finished after {} timesteps".format(t+1))
            break
```

Annotations for the code block:

- create a gym object (points to `env = gym.make('CartPole-v0')`)
- Initialize state (points to `observation = env.reset()`)
- figure (points to `env.render()`)
- Asking to GYM box (points to `env.step(action)`)

OpenAI Gym framework: CartPole game



State: 4 dimension vector

- [0]: cart position [-2.4 to 2.4]
- [1]: cart velocity [-Inf to Inf]
- [2]: pole angle [-42.8° to 41.8°]
- [3]: pole velocity at tip [-Inf to Inf]

[-0.00907215 0.01075882 -0.0133813 0.01925749]



```
import gym
env = gym.make('CartPole-v0')
for i_episode in range(20):
    observation = env.reset()
    for t in range(100):
        env.render()
        print(observation)
        action = env.action_space.sample()
        observation, reward, done, info = env.step(action)
        if done:
            print("Episode finished after {} timesteps".format(t+1))
            break
```

create a gym object

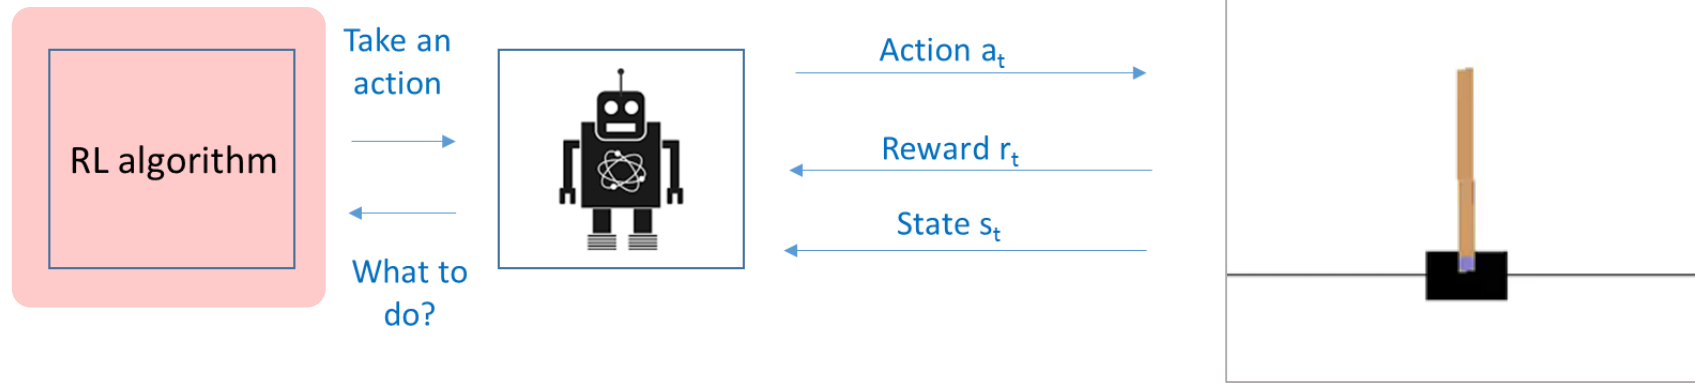
Initialize state

figure

Asking to GYM box

0:left, 1:right

Reinforcement Learning (RL) algorithm



- 1) Deep Q-Networks (DQN)
- 2) Policy Gradient (PG)
- 3) Actor Critic (AC)

Carpole game: basic information

<https://github.com/openai/gym/wiki/CartPole-v0>

❑ Action:

- 0: left
- 1: right

❑ State: 4 dimension vector

- [0]: cart position [-2.4 to 2.4]
- [1]: cart velocity [-Inf to Inf]
- [2]: pole angle [-42.8° to 41.8°]
- [3]: pole velocity at tip [-Inf to Inf]

❑ Initial state

- Random values between ± 0.05

❑ Reward:

- +1 each unit time if it is not fallen


❑ Episode Termination

- Pole Angle is more than $\pm 12^\circ$
- Cart Position is more than ± 2.4
- Episode length is greater than 200 (unit time)

❑ Training Termination

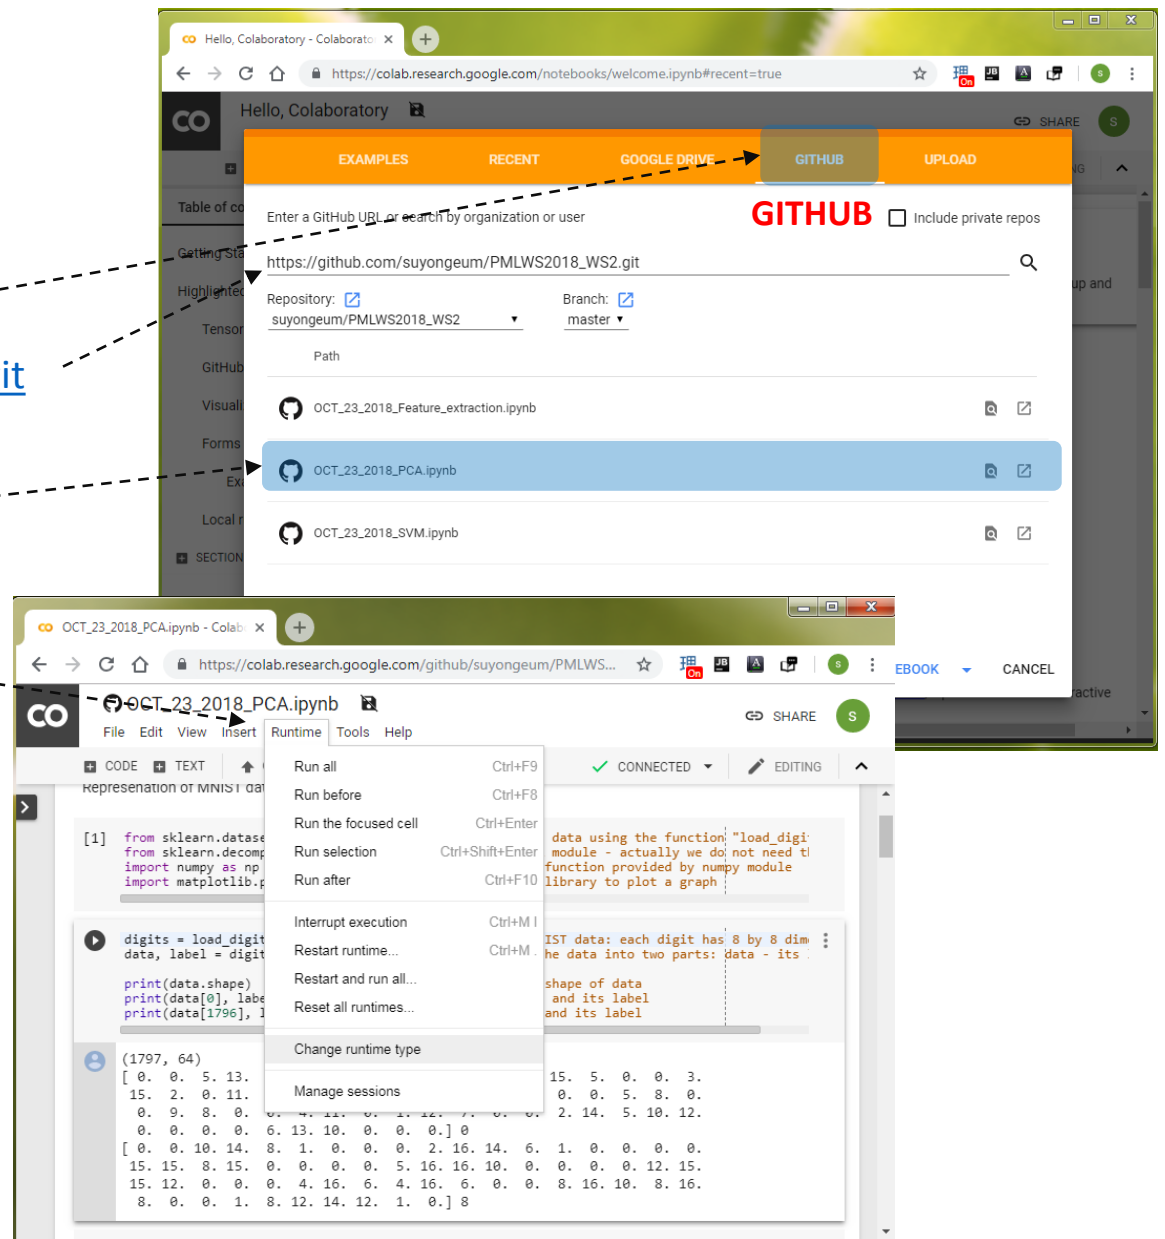
- Average reward is greater than or equal to 195 over 100 consecutive trials

```
# Get new state and reward from environment
next_state, reward, done, _ = env.step(action)
```



```
≡ 10 = {ndarray} [-0.1062376 -0.37924474 0.19689548 0.9208114 ]
```

- 1) Go to the Colab
 - <https://colab.research.google.com>
- 2) Select "GITHUB" and copy the link below into
 - https://github.com/suyongeuem/PMLWS2018_WS5.git
- 3) Select the notebook in the list
 - Dec_4_2018_cartpole.ipynb
- 4) Go to "Runtime" – "Change runtime type"
 - Python 3
 - GPU
- 5) Save it into your gdrive
 - "File" - "Save a copy in Drive ..."




Backup Slides

Cross Entropy


- Do you remember the Cross Entropy which is used to calculate the loss in CNN?
 - prediction: predicted label which is the output from the previous layer
 - e.g., [0.1, 0.2, 0.7]
 - label: true label, one-hot encoded
 - e.g., [0,0,1]

$$H(p, q) = -\sum_x p(x) \log q(\hat{x}) = -\log q(\hat{x}_{x \neq 0})$$

label prediction



$$-0 \cdot \log(0.1) - 0 \cdot \log(0.2) - 1 \cdot \log(0.7) = 0.375$$



$$-0 \cdot \log(0.3) - 0 \cdot \log(0.3) - 1 \cdot \log(0.4) = 0.916$$

Good one has small error