

② Deep Q-learning **DQN**

Published in 2013 - 2015 Nature

Approach: Q-learning \rightarrow what if q-values are represented by a **neural network**?

Bellman equations: Q^* is the **only** solution to that equation

$$Q^*(s,a) = \mathbb{E}_{r,s'} \left[r + \gamma \max_{a' \in A} Q^*(s',a') \right]$$

$\sim \Delta(s,a)$

$$\mathcal{L}(\theta) = \frac{1}{2} \left(\mathbb{E}_{(s,a,r,s')} \left[\underbrace{r + \gamma \max_{a' \in A} Q_{\theta'}(s',a')}_{\text{target model}} - \underbrace{Q_{\theta}(s,a)}_{\text{prediction model}} \right]^2 \right)$$

$\sim \sigma_{\theta}$

two occurrences of θ

[if $\mathcal{L}(\theta) = 0$ then Q_{θ} satisfies Bellman equation $\Rightarrow Q_{\theta} = Q^*$]

Key idea: use two networks!

↳ similarly to off-policy learning

↳ similarly to the maximisation bias

$$\mathcal{L}(\theta) = \frac{1}{2} \left(\mathbb{E}_{(s,a,r,s') \sim \sigma_{\theta}} \left[\underbrace{r + \gamma \max_{a' \in A} Q(s', a')}_{\text{fixed}} - Q(s, a) \right]^2 \right)$$

Implementation details:

- (prioritised) experience replay
- reward clipping

DQN algorithm

initialise two models : prediction model θ
target model θ'

$$\theta \rightarrow \sigma_{\theta}(s) = \arg \max_{a \in A} q_{\theta}(s, a)$$

Iterate:

First it: Batch sampling (B)

constant: batch size

First step: Batch sampling (ϵ)

Simulate the environment using ϵ . grady from θ
giving full trajectories

→ we break them down into B steps:

$$(s, a, r, s')$$

Second step: Update the networks

* at every iteration, update the prediction model:

compute $\hat{\mathcal{L}}(\theta)$

$$\mathcal{L}(\theta) = \frac{1}{2} \left(\mathbb{E}_{(s, a, r, s') \sim \theta} \left[\underbrace{r + \gamma \max_{a' \in A} Q(s', a')}_{\text{fixed}} - Q_{\theta}(s, a) \right] \right)^2$$

$\hat{\mathcal{L}}(\theta) = \frac{1}{2}$ average over steps from the batch:

$$(s, a, r, s') : r + \gamma \max_{a' \in A} Q(s', a') - Q_{\theta}(s, a)$$

if you have a neural network / ML model parametrised by θ , you have functions for

computing ∇_{θ} and for updating:

$$\hat{\mathcal{L}}(\theta) \xrightarrow{\text{gradient}} \nabla_{\theta} \hat{\mathcal{L}}(\theta) \xrightarrow{\text{optimisation step}} \theta_{\text{new}}$$

* every N iterations, update the target model:

$\Theta' \leftarrow \Theta$ ← parameters of the prediction model

↑ parameters of the target model

$N \approx 10$