

Artificial Neural Networks (Gerstner). Exercises for week 11

Deep Reinforcement Learning 2

Exercise 1. Uncorrelated mini-batches in A2C.

In the lecture you have seen a simple example of a single weight w that changes with temporally correlated updates Δw (red dots and red curve on slide 6). Reshuffling the updates in time led to a more stable learning dynamics (blue dots and blue curve). This example illustrates the effect of sampling iid from the replay buffer for off-policy methods like DQN. For on-policy methods, the proposed solution is to run multiple actors in parallel.

- Sketch a figure similar to the one on slide 6 for 4 parallel actors and 2 to 3 episodes per actor. Mark the starts of new episodes for each actor with vertical lines. Hint: the episodes can have different lengths.
- Draw in the same figure approximately the values $\frac{1}{4} \sum_{k=1}^4 \Delta w^{(k)}$ for all time points.
- Write a caption to the figure that explains, why the proposed solution of A2C helps to stabilize learning.

Exercise 2. Proximal Policy Optimization.

- In the derivation of Proximal Policy Optimization methods the ratio $r_{\theta'}(s_t, a_t) = \frac{\pi_{\theta'}(a_t; s_t)}{\pi_{\theta}(a_t; s_t)}$ appeared on the last line on slide 17. Convince yourself that the equality on the last line of slide 17 is correct by explicitly writing out the expectations in the same way as we did on slide 16.
Hint: Write something like $E_{s_t, a_t \sim p_{\theta'}, \pi_{\theta'}} [\sum_{t=0}^{\infty} \gamma^t A_{\theta}(s_t, a_t)] = \sum \dots = \sum \dots = E_{s_t, a_t \sim p_{\theta'}, \pi_{\theta} \dots}$
- Show that the summands in the loss function of PPO-CLIP can also be written in the form

$$\ell(r_{\theta'}) = \min(r_{\theta'} \gamma^t A_{\theta}, g(\epsilon, \gamma^t A_{\theta})), \quad \text{with } g(\epsilon, A) = \begin{cases} (1 + \epsilon)A & A \geq 0 \\ (1 - \epsilon)A & A < 0 \end{cases}$$

- Sketch $\ell(r)$ as a function of r in two figures: one where A_{θ} is positive and one where A_{θ} is negative.
- Write a caption to your figures that explains, why one can safely run a few steps of gradient ascent on $\hat{L}^{\text{CLIP}}(\theta')$ without risking that $\pi_{\theta'}$ would move too far away from π_{θ} .

Exercise 3. Deep Deterministic Policy Gradient.

- How many input and output neurons does the Q-network of DQN have, if the input consists of 100-dimensional vectors and there are 10 possible actions?
- How many input and output neurons does the Q-network of DDPG have, if the input consists of 100-dimensional vectors and the action space is 10 dimensional?
- Explain, why it would not be a good idea to use $\hat{Q}(s_{j+1}, a_{j+1})$ in line 7 of the DDPG algorithm (slide 24).
- Explain, why it would not be a good idea to use $\hat{Q}(s_{j+1}, \hat{\pi}(s_{j+1}) + \epsilon)$ in line 7 of the DDPG algorithm.

Exercise 4. Background Planning.

In this exercise we look again at the simple map of Europe on slide 29. You will run value iteration with goal Vienna. This means that we will keep $V(\text{V}) = 0$ all the time.

- Initialize all V- and Q-values to zero.
- Apply the update rule of value iteration (equation 1 on slide 29) to all cities in parallel. Hint 1: keep $V(\text{V}) = 0$ and don't worry, if you find e.g. $V(\text{Z}) = -2$. Hint 2: "In parallel" means, that you should assume $V(s') = 0$ when running this step for the first time and take the value that you obtained in the previous iteration otherwise.
- Repeat step b. until convergence.
- Convince yourself that value iteration found the optimal solution. Write down, how you convinced yourself that the optimal solution was found.