## Artificial Neural Networks (Gerstner). Exercises for week 8

## Deep Reinforcement Learning

## 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 $\Delta w$ (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.
a. 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.
b. Draw in the same figure approximately the values $\frac{1}{4} \sum_{k=1}^{4} \Delta w^{(k)}$ for all time points.
c. Write a caption to the figure that explains, why the proposed solution of A2C helps to stabilize learning.

## Exercise 2. Importance sampling.

Let us assume we would like to evaluate a policy $\pi(a \mid s)$, but we can only obtain episodes

$$
\left(S_{0}, A_{0}, R_{1}, S_{1}, \ldots, S_{T-1}, A_{T-1}, R_{T}, S_{T}\right)
$$

with policy $b(a \mid s)$. We will use importance weights $C_{t}$ to correct for the mismatch between the two policies, i.e. we will compute

$$
\tilde{V}_{\gamma}^{(T)}(b, s):=\mathbb{E}_{b}\left[\sum_{t=1}^{T} \gamma^{t-1} C_{t} R_{t} \mid S_{0}=s\right]
$$

where the expectation is taken over actions sampled from policy $b$. How should the importance weights $C_{t}$ be chosen to have $V_{\gamma}^{(T)}(\pi, s)=\tilde{V}_{\gamma}^{(T)}(b, s)$ ?
Hint: Importance weights are themselves random variable, i.e., they depends on ( $S_{0}, A_{0}, R_{1}, \ldots$ ).

## Exercise 3. Q-learning with function approximation



Figure 1: MDP in Exercise 3
Consider the MDP shown in Figure 1, with two states, two actions and deterministic rewards (where $T$ represents the terminal state). We want to learn the Q -values associated with the states using Q -learning, with discount factor $\gamma=1$.
a. (Tabular Case) The agent starts with all Q -values equal to 0 . We assume that the agent can store observed transitions in memory (similar to a replay buffer). The agent observes all 4 possible transitions, then updates the Q -values for $s_{2}$ by alternating between observations of ( $s_{2}, a_{1}, r_{21}$ ) and ( $s_{2}, a_{2}, r_{22}$ ) until learning converges. The agent then similarly alternates between observations of ( $s_{1}, a_{1}, r_{11}$ ) and $\left(s_{1}, a_{2}, r_{12}\right)$ until learning converges.
(i) What are the Q -values after convergence in $s_{2}$, and finally after convergence in $s_{1}$ ?
(ii) Do the Q -values after each stage result in the optimal policy?
b. (Function Approximation) Now assume that the states are given to us with the vector-based observations shown in Figure 2 left. We will learn the Q-values using the linear network shown in Figure 2 right.
As before, assume a replay-buffer-style learning where the agent learns the weights after observing all transitions. Start with $w_{11}=w_{12}=w_{13}=w_{21}=w_{22}=w_{23}=0$.
(i) What will the converged weights be after alternating between the two possible $s_{2}$ observations? Hint: Note that certain weights will always be updated in exactly the same way and should, therefore, converge to the same value.
(ii) After $s_{2}$ convergence, what is the policy in $s_{1}$ ? How does this differ from the tabular case after $s_{2}$ convergence, and why?
(iii) What weights would result in the correct Q -value predictions for all ( $s, a$ ) pairs? Are they unique?
(iv) How can an arbitrary tabular Q-learning problem be represented using a simple linear neural network like the one shown on the right?
Hint: consider how the input space could be represented such that semi-gradient descent results in each weight converging exactly to $Q(s, a)$ for some $(s, a)$ pair.

$$
\begin{aligned}
& s_{1}=\left(\begin{array}{l}
x_{1} \\
x_{2} \\
x_{3}
\end{array}\right)=\left(\begin{array}{l}
1 \\
1 \\
0
\end{array}\right) \\
& s_{2}=\left(\begin{array}{l}
x_{1} \\
x_{2} \\
x_{3}
\end{array}\right)=\left(\begin{array}{l}
0 \\
1 \\
1
\end{array}\right)
\end{aligned}
$$



Figure 2: Neural network for function approximation in Exercise 3

## Exercise 4. Proximal Policy Optimization.

a. In the derivation of Proximal Policy Optimization methods the ratio $r_{\theta^{\prime}}\left(s_{t}, a_{t}\right)=\frac{\pi_{\theta^{\prime}}\left(a_{t} ; s_{t}\right)}{\pi_{\theta}\left(a_{t} ; s_{t}\right)}$ appeared on the last line on slide 22 . Convince yourself that the equality on the last line of slide 22 is correct by explicitly writing out the expectations in the same way as we did on slide 21.
Hint: Write something like $E_{s_{t}, a_{t} \sim p_{\theta^{\prime}}, \pi_{\theta^{\prime}}}\left[\sum_{t=0}^{\infty} \gamma^{t} A_{\theta}\left(s_{t}, a_{t}\right)\right]=\sum \ldots=\sum \ldots=E_{s_{t}, a_{t} \sim p_{\theta^{\prime}}, \pi_{\theta}} \ldots$
b. Show that the summands in the loss function of PPO-CLIP can also be written in the form

$$
\ell\left(r_{\theta^{\prime}}\right)=\min \left(r_{\theta^{\prime}} \gamma^{t} A_{\theta}, g\left(\epsilon, \gamma^{t} A_{\theta}\right)\right), \quad \text { with } g(\epsilon, A)=\left\{\begin{array}{cc}
(1+\epsilon) A & A \geq 0 \\
(1-\epsilon) A & A<0
\end{array}\right.
$$

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

## Exercise 5. Deep Deterministic Policy Gradient.

a. How many input and output neurons does the Q-network of DQN have, if the input consists of 100dimensional vectors and there are 10 possible actions?
b. How many input and output neurons does the Q-network of DDPG have, if the input consists of 100dimensional vectors and the action space is 10 dimensional?
c. Explain, why it would not be a good idea to use $\hat{Q}\left(s_{j+1}, a_{j+1}\right)$ in line 7 of the DDPG algorithm (slide 20).
d. Explain, why it would not be a good idea to use $\hat{Q}\left(s_{j+1}, \hat{\pi}\left(s_{j+1}\right)+\epsilon\right)$ in line 7 of the DDPG algorithm.

