## Artificial Neural Networks (Gerstner). Exercises for week 1

## Reinforcement Learning: Basics

## Exercise 1. Iterative update ${ }^{1}$

We consider an empirical evaluation of $Q(s, a)$ by averaging the rewards for action $a$ over the first $k$ trials:

$$
Q_{k}=\frac{1}{k} \sum_{i=1}^{k} r_{i} .
$$

We now include an additional trial and average over all $k+1$ trials.
a. Show that this procedure leads to an iterative update rule of the form

$$
\Delta Q_{k}=\eta_{k}\left(r_{k}-Q_{k-1}\right),
$$

(assuming $Q_{0}=0$ ).
b. What is the value of $\eta_{k}$ ?
c. Give an intuitive explanation of the update rule.

Hint: Think of the following: If the actual reward is larger than my estimate, then I should ...

## Exercise 2. Greedy policy and the two-armed bandit

In the " 2 -armed bandit" problem, one has to choose one of 2 actions. Assume action $a_{1}$ yields a reward of $r=1$ with probability $p=0.25$ and 0 otherwise. If you take action $a_{2}$, you will receive a reward of $r=0.4$ with probability $p=0.75$ and 0 otherwise. The " 2 -armed bandit" game is played several times and Q values are updated using the update rule $\Delta Q(s, a)=\eta\left[r_{t}-Q(s, a)\right]$.
a. Assume that you initialize all Q values at zero. You first try both actions: in trial 1 you choose $a_{1}$ and get $r=1$; in trial 2 you choose $a_{2}$ and get $r=0.4$. Update your Q values ( $\eta=0.2$ ).
b. In trials 3 to 5 , you play greedy and always choose the action which looks best (i.e., has the highest Q -value). Which action has the higher Q -value after trial 5? (Assume that the actual reward is $r=0$ in trials 3-5.)
c. Calculate the expected reward for both actions. Which one is the best?
d. Initialize both $Q$-values at 2 (optimistic). Assume that, as in the first part, in the first two trials you get for both actions the reward. Update your Q values once with $\eta=0.2$. Suppose now that in the following rounds, in order to explore well, you choose actions $a_{1}$ and $a_{2}$ alternatingly and update the Q -values with a very small learning rate ( $\eta=0.001$ ). How many rounds (one round $=$ two trials $=$ one trial with each action) does it take on average, until the maximal Q -value also reflects the best action?
Hint: For $\eta \ll 1$ we can approximate the actual returns $r_{t}$ with their expectations $E[r]$.

## Exercise 3. Batch vesrsus online learning rules: Recap ${ }^{2}$

We define the mean squared error in a dataset with $P$ data points as

$$
\begin{equation*}
E^{\mathrm{MSE}}(\boldsymbol{w})=\frac{1}{2} \frac{1}{P} \sum_{\mu}\left(t^{\mu}-\hat{y}^{\mu}\right)^{2} \tag{1}
\end{equation*}
$$

[^0]where the output is
\[

$$
\begin{equation*}
\hat{y}^{\mu}=g\left(a^{\mu}\right)=g\left(\boldsymbol{w}^{T} \boldsymbol{x}^{\mu}\right)=g\left(\sum_{k} w_{k} x_{k}^{\mu}\right) \tag{2}
\end{equation*}
$$

\]

and the input is the $\boldsymbol{x}^{\mu}$ with components $x_{1}^{\mu} \ldots x_{d}^{\mu}$.
a. Calculate the update of weight $w_{j}$ by gradient descent (batch rule)

$$
\begin{equation*}
\Delta w_{j}=-\eta \frac{d E}{d w_{j}} \tag{3}
\end{equation*}
$$

Hint: Apply chain rule
b. Rewrite the formula by taking one data point at a time (stochastic gradient descent). What is the difference to the batch rule?
c. Rewrite your result in $b$ in vector notation (hint: use the weight vector $\boldsymbol{w}$ and the input vector $\left.\boldsymbol{x}^{\mu}\right)$. Show that the update after application of data point $\mu$ can be written as

$$
\Delta \boldsymbol{w}=\eta \delta(\mu) \boldsymbol{x}^{\mu}
$$

where $\delta(\mu)$ is a scalar number that depends on $\mu$. Express $\delta(\mu)$ in terms of $t^{\mu}, \hat{y}^{\mu}, g^{\prime}$.

## Exercise 4. Geometric interpretation of an artificial neuron: Recap

Consider the single-neuron function in 2-D with

$$
\begin{equation*}
y=g\left(\boldsymbol{x}^{T} \boldsymbol{w}\right) \tag{4}
\end{equation*}
$$

where $g$ is a strictily increasing activation function, $\boldsymbol{x}=\left(x_{1}, x_{2},-1\right) \in \mathbb{R}^{2+1}$ is the extended 2dimensional input (i.e., the threshold/bias value has been integrated as an extra input $x_{3}=-1$ ), and $\boldsymbol{w}=\left(w_{1}, w_{2}, w_{3}\right) \in \mathbb{R}^{3}$ is the weight vector. The hyperplane $\boldsymbol{x}^{T} \boldsymbol{w}=0$ describes the boundary between where the neuron is on, i.e., $\boldsymbol{x}^{T} \boldsymbol{w}>0$, and where it is off, i.e., $\boldsymbol{x}^{T} \boldsymbol{w}<0$. Consider this hyperplane in the $2-\mathrm{D}$ space of $\left(x_{1}, x_{2}\right)$ and answer the following questions:
a. The hyperplane is a line in $2-\mathrm{D}$. What is the slope of this line as a function of $w_{1}, w_{2}$, and $w_{3}$ ? Where does the line intersect with the $y$-axis and where with the $x$-axis?
b. Is it possible to have two weight vectors $\boldsymbol{w}$ and $\boldsymbol{w}^{\prime}$ such that $\boldsymbol{w} \neq \boldsymbol{w}^{\prime}$ but $\boldsymbol{x}^{T} \boldsymbol{w}=0$ and $\boldsymbol{x}^{T} \boldsymbol{w}^{\prime}=0$ describe the same hyperplane? If yes, what conditions $\boldsymbol{w}$ and $\boldsymbol{w}^{\prime}$ must meet?
c. For the general case of $\boldsymbol{x}=\left(x_{1}, \ldots, x_{N},-1\right) \in \mathbb{R}^{N+1}$, what is the distance of the hyperplane $\boldsymbol{x}^{T} \boldsymbol{w}=0$ from the origin in $\mathbb{R}^{N}$ ? Where does the hyperplane intersect with the $x_{n}$-axis for $n \in\{1, \ldots, N\}$ ?
d. Use the online learning rule you derived in Exercise 3c and describe, in words, how the separating hyperplane in $\mathbb{R}^{N}$ changes after each update. Make sure you consider the effects of both changing bias/threshold on one side and changing weight parameter $\boldsymbol{w}$ on the other side.


[^0]:    ${ }^{1}$ The result of Exercise 1 will be used in the second lecture of week 1 .
    ${ }^{2}$ The result of Exercise 3 will be used in the second lecture of week 1 .

