## Artificial Neural Networks (Gerstner). Exercises for week 3

## Markov Decision Processes

## Exercise 1. Optimal policies for finite horizon.

Create a Markov Decision Process where the optimal horizon- $T$ policy depends on the time step, i.e. there is at least one state $s$ and one pair of timesteps $t$ and $t^{\prime}$ such that $\pi^{(t)}(a \mid s) \neq \pi^{\left(t^{\prime}\right)}(a \mid s)$.
Hint: You can choose $T=2$ for simplicity.

## Exercise 2. Shortest path search.

Let $\mathcal{S}=\left\{s_{1}, s_{2}, s_{3}, \ldots\right\}$ denote a set of vertices (think of cities on a map) and let the vertices be connected by some edges $e_{s_{i}, s_{j}} \in(0, \infty]$ (think of distances between cities), where $e_{s_{i}, s_{j}}=\infty$ indicates that there is no direct connection between $s_{i}$ and $s_{j}$. Dijkstra's algorithm for finding the shortest paths to some goal vertex $g$ can be written in the following way (we show the lenght of the shortest path from vertex $s$ to $g$ by $V(s))$ :

- For each vertex $s \in \mathcal{S}$, initialize all distances from $g$ by $V(s) \leftarrow \infty$.
- Initialize the distance of $g$ from itself by $V(g) \leftarrow 0$.
- Define and initialize $\tilde{\mathcal{S}} \leftarrow \mathcal{S}$.
- While $\tilde{\mathcal{S}}$ is not empty
$-s_{i} \leftarrow \arg \min _{s \in \tilde{\mathcal{S}}} V(s)$
- Remove $s_{i}$ from $\tilde{\mathcal{S}}$
- For each neighbor $s_{j}$ of $s_{i}$ still in $\tilde{\mathcal{S}}: V\left(s_{j}\right) \leftarrow \min \left(V\left(s_{j}\right), V\left(s_{i}\right)+e_{s_{i}, s_{j}}\right)$.
- Return $V(s)$ for all $s \in \mathcal{S}$.

The output $V(s)$ of Dijkstra's algorithm is equal to the lenght of the shortest path from $s$ to $g$. In this exercise, we formulate the problem of finding the shortest path as a dynamic programming problem.
a. What is the equivalent Markov Decision Process for the problem of finding the shortest paths to some goal state?
Hint: Define the goal state as an absorbing state and describe the properties of $r_{s}^{a}$ and $p_{s_{i} \rightarrow s_{j}}^{a}$.
b. Compare the value iteration algorithm on the MDP of part a with Dijkstra's algorithm.

## Exercise 3. Bellman operator.

Proof that the Bellman operator is a contraction.
Hint: Show the contraction with the infinity norm, i.e.

$$
\left\|T_{\gamma}[X]-T_{\gamma}[Y]\right\|_{\infty}=\max _{s}\left|T_{\gamma}[X]_{s}-T_{\gamma}[Y]_{s}\right| \leq \gamma\|X-Y\|_{\infty},
$$

where the last inequality is to be proven. You can use the notation $Q_{s a}^{X}=r_{s}^{a}+\gamma \sum_{s^{\prime} \in \mathcal{S}} p_{s \rightarrow s^{\prime}}^{a} X_{s^{\prime}}$ and the facts that $\left|\max _{a} Q_{s a}^{X}-\max _{a^{\prime}} Q_{s a^{\prime}}^{Y}\right| \leq \max _{a}\left|Q_{s a}^{X}-Q_{s a}^{Y}\right|$ and $\sum_{s^{\prime} \in \mathcal{S}} p_{s \rightarrow s^{\prime}}^{a}=1$.

## Exercise 4. Coding exercise: Value and policy iteration.

Implement value and policy iteration in python to solve the MDP from Example 1 from the lecture.

