

Last Name .....

First Name.....

---

## Artificial Neural Networks: Exam

### 29th of June 2023

---

- Keep your bag next to your chair, but do not open it during the exam.
- Write your name in legible letters on top of this page.
- The exam lasts 180 min. (The estimated times per section add up to 160 min)
- Write **all** your answers in a legible way on the exam (no extra sheets).
- No documentation is allowed (no textbook, no slides), except **one page A5 of handwritten notes, doublesided**.
- No calculator is allowed.
- Have your student card displayed in front of you on your desk.
- **Check that your exam has 10 pages.**

Evaluation:

1. .... / 5 pts (Section 1 - estimated time 20 min)
2. .... / 7 pts (Section 2 - estimated time 35 min)
3. .... / 8 pts (Section 3 - estimated time 35 min)
4. .... / 12 pts (Section 4 - estimated time 70 min)

---

**Total: .... / 32 pts**

This page remains empty. You can use it as free space for your calculations, do not use to write down answers.

## Definitions and notations

RL stands for Reinforcement Learning.

The symbol  $\eta$  is reserved for the learning rate.

Bold face symbols refer to vectors, normal face to a single component or a single input/output. Unless noted otherwise, the input is  $N$ -dimensional:  $\mathbf{x}^u \in R^N$

In the context of reinforcement learning, the symbol  $a$  refers to an action; the symbols  $r$  and  $R$  to a reward; the symbol  $s$  to a discrete state; and the symbol  $\gamma$  to a discount rate.

If the state space is continuous then states are also written as  $\mathbf{x}$ .

If algorithms are implemented as neural networks, the parameters are the weights  $w_{ij}$  from neuron  $j$  to neuron  $i$ . In a multilayer network an upper index ( $k$ ) refers to the layer, so that we write  $w_{ij}^{(k)}$ . Hence  $w_{nm}^{(l)}$  refers to the connection from of neuron  $m$  in layer  $l - 1$  to neuron  $n$  in layer  $l$ .

In many RL applications, there is a single goal state (terminal state). An action sequence that ends at the goal is called an episode.

## How to give answers

All sections involve short calculations. **Please write the answers in the space provided for that purpose.**

We also provide some free space for calculations. We will not look at these parts for grading. You can ask for extra scratch paper. We will not look at the scratch paper.

**1 3-step SARSA (5 points). Estimated time: 20 minutes**

We approximate Q-values  $Q(\mathbf{x}, a)$  by function approximation in a deep neural network with weights  $w_{ij}^{(k)}$  where  $i = 1, 2, \dots$  and  $j = 1, 2, \dots$  are arbitrary neuron indices and  $k$  is the layer-index. The input  $\mathbf{x}$  is continuous.

(i) Write down a reasonable loss function resulting from the consistency condition of the Bellman equation for the case of **3-step SARSA**. Consider both the BATCH version of the loss and the ONLINE version of the loss (the ONLINE version should contain only a minimal set of time steps necessary to implement SARSA updates). Choose an appropriate and understandable notation.

BATCH:

Loss  $L_{bt} = \dots\dots\dots$

ONLINE:

Loss  $L_{on} = \dots\dots\dots$

number of points: ...../ 2

(ii) Derive an online update rule for the specific weight  $w_{nm}^{(k)}$  starting from the ONLINE Loss  $L_{on}$  in (i) using **semigradient** (we write  $\Delta_{sg}$  and  $\Delta_{fg}$  for updates with semigradient and full gradient, respectively)

$\Delta_{sg}w_{nm}^{(k)} = \dots\dots\dots$

number of points: ...../ 2

(iii) Write down the **DIFFERENCE** between the online update rule with *semi-gradient*  $\Delta_{sg}w_{nm}^{(k)}$  on the ONLINE loss  $L_{on}$  in (i) and that with the *full gradient*  $\Delta_{fg}w_{nm}^{(k)}$  (standard stochastic gradient) on the ONLINE loss  $L_{on}$  in (i).

DIFFERENCE

$\Delta_{sg}w_{nm}^{(k)} - \Delta_{fg}w_{nm}^{(k)} =$   
 .....

number of points: ...../ 1

**Space for your calculations, not for answers**

**2 Policy Gradient for 1-step horizon: Contextual Bandit Problem (7 points). Estimated time 35 min**

The agent (bandit) can choose amongst  $N$  actions  $a = 1, 2, \dots, N$ . Inputs are denoted by  $\mathbf{x}$ . Action  $a = i$  (with  $1 \leq i \leq N$ ) is chosen with probability

$$\pi(a = i | \mathbf{x}; \mathbf{w}_1, \dots, \mathbf{w}_N) = \frac{[\sum_{k=1}^{10} w_{ik} y_k]^\beta}{\sum_j [\sum_{k=1}^{10} w_{jk} y_k]^\beta}, \quad (1)$$

where  $\beta \in \{2, 4, 6, \dots\}$  is a metaparameter,  $\mathbf{w}_i = (w_{i1}, \dots, w_{i10})$  is the weight vector for action  $a = i$ , and  $y_k = f(\mathbf{x} - \mathbf{c}_k)$  with  $1 \leq k \leq 10$  is a set of ten basis functions. The agent reacts to an input  $\mathbf{x}$  with an action  $a$  and receives a deterministic reward  $r(\mathbf{x}, a)$ . At each time step a new input  $\mathbf{x}$  is drawn from a stationary distribution  $P(\mathbf{x})$ .

(i) Write down the expected reward  $E[r(\mathbf{x}, a) | \mathbf{w}_1, \dots, \mathbf{w}_N]$  for the policy with parameters  $\mathbf{w}_1, \dots, \mathbf{w}_N$ :

$$E[r(\mathbf{x}, a) | \mathbf{w}_1, \dots, \mathbf{w}_N] =$$

.....

number of points: ...../ [1]

(ii) What is the stochastic online version of the gradient  $E[r(\mathbf{x}, a) | \mathbf{w}_1, \dots, \mathbf{w}_N]$  with respect to a weight  $w_{nm}$ ? Write down the result.

.....

number of points: ...../ [2]

(iii) Show that in expectation your online version in (ii) leads back to correct gradient of  $E[r(\mathbf{x}, a) | \mathbf{w}_1, \dots, \mathbf{w}_N]$

.....

.....

.....

number of points: ...../ [2]

(iv) Take your result from (ii), separate a common factor  $\eta\beta/(\sum_k w_{nk} y_k)$  from the remaining terms, and use Eq. (1) to find an elegant and compact formula for the update of the weight  $w_{nm}$ :

$$\Delta w_{nm} = \dots\dots\dots$$

number of points: ...../ [2]

**3 1-step and 2-step Q-learning in the tabular setting (8 points). Estimated time: 35 min**

In class we have discussed the SARSA algorithm and shown that, if the Q-values have converged, the resulting set of Q-values solves in expectation the Bellman equation. Here the task is to do a related calculation for Q-learning.

The agent uses an epsilon-greedy policy. The update step is

$$\Delta Q(s_t, a_t) = \eta [r_t + \gamma \max_{a'} Q(s_{t+1}, a') - Q(s_t, a_t)] \quad (2)$$

The update step is applied after *every* action.

(i) Assume that the Q-learning algorithm has converged in expectation, i.e., the expectation of the update step is  $E[\Delta Q(s_t, a_t)] = 0$ . Show that  $E[\Delta Q(s, a)] = 0$  for all pairs  $(s, a)$  implies that the set of Q-values solves the optimal Bellman equation.

- (1) .....
- (2) .....
- (3) .....
- (4) .....

Line (4) is the final result. For the transitions from lines (1) to (2) to (3) to (4) I used the following identities or arguments:

- (1) to (2) .....
- (2) to (3) .....
- (3) to (4) .....

number of points: ...../ [2]

(ii) So far we have looked at the standard version of Q-learning which we call Variant A<sub>1</sub> in the following. Your friend Bob proposes an alternative (Variant B) of the algorithm as follows: As before we run an epsilon-greedy policy, but to be sure that we remain optimal, we apply the update step (given above in Eq. (2)) only if the *best* action is chosen at time  $t$ , i.e.,  $a_t = a^*(s_t) = \max_{a'} Q(s_t, a')$ . Your friend Alice says that convergence of Q-values  $E[\Delta Q(s, a)] = 0$  for all pairs  $(s, a)$  in Variant B **does not guarantee** that the set of Q-values is consistent with the Bellman equation.

Is Alice right? Give a mathematical argument or an intuitive interpretation.

.....  
.....

number of points: ...../ [2]

(iii) Now we repeat the same arguments for 2-step Q-learning. Again we consider two variants.

In variant  $A_2$  the update step is applied after *every* action (for all time steps  $t > 2$ ). Your friend Claire proposes an alternative (Variant C) where the update rule is applied only if the epsilon-greedy algorithm chooses the greedy action for the intermediate step of the algorithm, i.e.  $a_{t+1} = a^*(s_{t+1}) = \max_{a'} Q(s_{t+1}, a')$ .

Your friend Dominique claims that (at least) one of the Variants  $A_2$  or C is consistent with the Bellman equation. To check whether Dominique is correct, let us assume that the 2-step Q-learning algorithm has converged in expectation, i.e., the expectation of the update step is  $E[\Delta Q(s_t, a_t)] = 0$ .

Choose either variant  $A_2$  or C (whichever works best for your arguments) and show that  $E[\Delta Q(s_t, a_t)] = 0$  implies that the set of Q-values solves the two-step Bellman equation. Write down the two-step algorithm, then three steps of the argument and then your final result:

.....  
.....  
.....  
.....  
.....

number of points: ...../ [3]

(iv) Would the mathematical calculation also work for the other variant? The answer (yes/no) needs to be supported by a mathematical argument or an intuitive reason.

.....  
.....

number of points: ...../ [1]

**4 Importance Sampling in an n-step actor-critic model: V-Trace approach (12 points). Estimated time: 70 min.**

*Background information, not necessary to solve the exercise:* The V-Trace Actor-Critic Method is an important part of the popular IMPALA agent. It is an *off-policy* modification of an n-step actor-critic policy gradient method. In this section we analyze aspects of this practically relevant algorithm.

*Problem setting:* We consider trajectories as sequences of states  $\mathbf{x}_t, \mathbf{x}_{t+1}, \mathbf{x}_{t+2}, \dots$  and actions  $a_t, a_{t+1}, a_{t+2}, \dots$  that are generated by an actor following some policy  $\mu(a|\mathbf{x})$ . We are interested in the V-values  $V(\mathbf{x})$  that correspond to another policy  $\pi(a|\mathbf{x})$ . The state at time  $t$  is denoted by  $\mathbf{x}_t$ ; the state at time  $s$  is denoted by  $\mathbf{x}_s$ .

The critic is optimized by semi-gradient on a loss function  $L_{on} = (V(\mathbf{x}_s) - v_s)^2$  that compares, at each time point  $s$ , the V-value  $V(\mathbf{x}_s)$  in state  $\mathbf{x}_s$  with a target  $v_s$ . The actor is optimized by policy gradient.

*V-trace target:* We define the  $n$ -step 'V-Trace target' for  $V(\mathbf{x}_s)$  as

$$v_s := V(\mathbf{x}_s) + \sum_{t=s}^{s+n-1} \gamma^{t-s} \left( \prod_{i=s}^{t-1} c_i \right) \delta_t \tag{3}$$

where

$$\delta_t := [r_t + \gamma V(\mathbf{x}_{t+1}) - V(\mathbf{x}_t)] \rho_t \tag{4}$$

is a temporal difference and  $\rho_t := \min\left(\bar{\rho}, \frac{\pi(a_t|\mathbf{x}_t)}{\mu(a_t|\mathbf{x}_t)}\right)$  and  $c_i := \min\left(\bar{c}, \frac{\pi(a_i|\mathbf{x}_i)}{\mu(a_i|\mathbf{x}_i)}\right)$  are truncated importance sampling weights with truncation levels  $\bar{\rho}$  and  $\bar{c}$ .

Note: To simplify the writing, we make use of the notation  $\prod_{i=s}^{t-1} c_i = 1$  for  $s = t$ .

(i) We start with the on-policy case, assuming  $\bar{\rho} = \bar{c} = 1$ . Rewrite the V-Trace target  $v_s$  defined by Eq. 3 in terms of  $V(\mathbf{x}_{s+n})$  and intermediate rewards  $r_s, r_{s+1}, \dots, r_{s+n}$ .

*Hint1:* After the rewrite,  $V(\mathbf{x}_s)$  should disappear on the right-hand-side of Eq. 3.

*Hint2:* on-policy means  $\mu = \pi$ .

$v_s = \dots\dots\dots$

number of points: ...../ 3

**Space for your calculations, not for answers**



(ii) We turn to the critic and study the off-policy scenario. Hence the policy  $\pi \neq \mu$ .

Expectation values will be denoted by  $E_\mu$  if the expectation is over trajectories generated by the policy  $\mu$  and  $E_\pi$  if the expectation is over trajectories generated by the policy  $\pi$ .

For the calculation we assume  $\bar{\rho} = \bar{c} > \max_t \frac{\pi(a_t|\mathbf{x}_t)}{\mu(a_t|\mathbf{x}_t)}$ .

Given the current state  $\mathbf{x}_s$ , relate the Expectation  $E_\mu$  of V-Trace target  $v_s$  from Eq. 3 to the Expectation  $E_\pi$  of V-values and, possibly, rewards.

Show the main steps of your mathematical calculation and write the final result in the last line.

Hint: You start with  $E_\mu[v_s|\mathbf{x}_s]$  on the left-hand side; and end with an expectation  $= E_\pi[\dots]$  over an appropriate expression on the right-hand side.

$$E_\mu[v_s|\mathbf{x}_s] = E_\mu[ \dots ]$$
$$= \dots$$

$$E_\mu[v_s|\mathbf{x}_s] = E_\pi[ \dots ]$$

number of points: ...../ 4

**Space for your calculations, not for answers**

**Exam continues on next page**

(iii) We now turn to the actor. We consider the term

$$A := \rho_s \nabla_w \log \pi_w(a_s | \mathbf{x}_s) (r_s + \gamma v_{s+1} - V(\mathbf{x}_s)) \tag{5}$$

and study its expectation  $E_\mu[A | \mathbf{x}_s]$ . We assume  $\bar{\rho} = \bar{c} > \max_t \frac{\pi(a_t | \mathbf{x}_t)}{\mu(a_t | \mathbf{x}_t)}$ .

Show that  $E_\mu[A | \mathbf{x}_s]$  is equal to the expectation  $E_\pi$  of the policy-gradient rule of the **on-policy model**, i.e., of an agent with policy  $\pi$ .

Show the main steps of your mathematical arguments:

(1)  $E_\mu[A | \mathbf{x}_s] = \dots\dots\dots$

(2)  $\quad = \dots\dots\dots$

(3)  $\quad = \dots\dots\dots$

(4)  $\quad = \dots\dots\dots$

Line 4, is the final result. For the transitions from line 1 to 2 to 3 to 4, I used the following identities:

(1) to (2)  $\dots\dots\dots$

(2) to (3)  $\dots\dots\dots$

(3) to (4)  $\dots\dots\dots$

number of points: ...../ 4

(iv) Why are the findings in parts (ii) and (iii) important for practical applications?

.....  
 .....  
 .....  
 .....

number of points: ...../ 1

**Space for your calculations, not for answers**