

## Intelligent Systems, 2022/23, written exam, 28 August 2023

All questions count equally. Literature, electronic and communication devices are not allowed. It is allowed to use one sheet of A4 format paper with notes. You can write your answers in English or Slovene. Duration: 90 minutes.

A possibility to see gradings of your exams will take place on Thursday, 31 August 2023, at 10:00 in the office of Prof Robnik Šikonja (2<sup>nd</sup> floor, room 2.06).

1. A student is training a deep neural network to classify images. The network has several layers with a large number of neurons. The student is using the sigmoid activation function in the hidden layers. After training the network for several epochs, the student found that the model's performance is not improving and suspected the problem of vanishing gradients.

To alleviate the problem, the student considered four techniques, listed below. Describe if and how would the considered techniques alleviate the vanishing gradients problem:

- a) ReLU activation function,
- b) skip connections,
- c) an optimizer with adaptive learning rate,
- d) dropout.

2. A medical research team is trying to develop a machine learning model to predict if a patient has a rare disease. The dataset has only 107 positive cases of the disease and 12,155 negative cases. The team used accuracy as the evaluation metric for their model.

- a) What are the problems with using accuracy as an evaluation metric for this dataset?

For the evaluation metrics given below explain if and why they would be more suitable for this type of dataset.

- b) sensitivity and specificity,
- c) precision, recall, and F1 metric,
- d) AUC.

3. A student is training a regression model to predict the price of a house based on its characteristics (40 attributes) and location (20 attributes). The student used feature selection before training a linear regression model and noticed that the model has high bias and is underfitting the data.

- a) What is the concept of bias in machine learning and how does it relate to underfitting?
- b) Explain two approaches how could the student reduce underfitting.
- c) Give an example of how the student could use regularization to improve the linear regression model by adjusting the regularization parameter lambda.

4. A student is working on a multiclass classification task, and has chosen to use a support vector machine classifier using one-versus-all strategy. After evaluating the model on a test set, the student found that the performance is not as good as expected.

- a) What are problems that the student may face when using SVM for multiclass classification?
- b) How can the student address problems related to class imbalance by using one-versus-one strategy?
- c) How can the student address problems related to class imbalance by using cost-sensitive learning?
- d) How can the student address problems related to class imbalance by using class balancing?

5. A student is working on developing an AI agent to play a video game. The student has chosen to use Q-learning as the reinforcement learning algorithm.

- a) What is the Q-learning algorithm and how does it work?
- b) Propose a set of states and actions, and propose a sensible reward for this problem.
- c) What are some of the challenges the student may face when implementing Q-learning for this task?